



**NAVAL  
POSTGRADUATE  
SCHOOL**

**MONTEREY, CALIFORNIA**

**THESIS**

**VIRTUALIZATION OF AEGIS: A STUDY OF THE  
FEASIBILITY OF APPLYING OPEN ARCHITECTURE  
TECHNOLOGY TO THE SURFACE NAVY'S MOST  
COMPLEX AUTOMATED WEAPON SYSTEM**

by

Erik S. Roberts

September 2011

Thesis Co-Advisors:

Man-Tak Shing  
Albert Barreto III

**Approved for public release; distribution is unlimited**

THIS PAGE INTENTIONALLY LEFT BLANK

**REPORT DOCUMENTATION PAGE**
*Form Approved OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.

<b>1. AGENCY USE ONLY (Leave blank)</b>		<b>2. REPORT DATE</b> September 2011	<b>3. REPORT TYPE AND DATES COVERED</b> Master's Thesis
<b>4. TITLE AND SUBTITLE</b> Virtualization of AEGIS: A Study of the Feasibility of Applying Open Architecture Technology to the Surface Navy's Most Complex Automated Weapon System		<b>5. FUNDING NUMBERS</b>	
<b>6. AUTHOR(S)</b> Erik S. Roberts			
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Naval Postgraduate School Monterey, CA 93943-5000		<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>	
<b>9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> N/A		<b>10. SPONSORING/MONITORING AGENCY REPORT NUMBER</b>	
<b>11. SUPPLEMENTARY NOTES</b> The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number <u>N/A</u> .			
<b>12a. DISTRIBUTION / AVAILABILITY STATEMENT</b> Approved for public release; distribution is unlimited		<b>12b. DISTRIBUTION CODE</b> A	
<b>13. ABSTRACT (maximum 200 words)</b> Rising costs of proprietary equipment in legacy electronic applications are increasingly drawing resources from vital programs. Growing interest in evaluating Open Architecture technology to replace closed systems is evidenced by the number of recent publications on the subject. Researchers have approached this topic from various angles, including lifecycle management, risk simulation, total cost of ownership, and knowledge-value added measures.  This exploratory study uses open architecture hardware employing virtualization technology to test the feasibility of replacing legacy components of military systems. Virtualization has the potential to provide significant cost savings in terms of procurement, daily operation, and maintenance. Additionally, virtualization provides functional benefits such as load-balancing, greater processor utilization and storage flexibility, streamlined scalability, and simplified disaster recovery strategies.  This thesis is original research in the form of a proof-of-concept study. It details performance results of a locally-constructed test platform, designed to simulate a portion of the U.S. Navy's AEGIS Weapon System. The scope of this work is to test the viability of using commodity-based hardware to achieve performance levels equal to, or greater than, current proprietary systems. Value-Added metrics are applied through cost comparisons between the test platform and typical AEGIS systems. While this study specifically targets AEGIS, the results can be generalized to non-military applications.			
<b>14. SUBJECT TERMS</b> AEGIS; Virtualization; Open Architecture			<b>15. NUMBER OF PAGES</b> 73
<b>16. PRICE CODE</b>			
<b>17. SECURITY CLASSIFICATION OF REPORT</b> Unclassified	<b>18. SECURITY CLASSIFICATION OF THIS PAGE</b> Unclassified	<b>19. SECURITY CLASSIFICATION OF ABSTRACT</b> Unclassified	<b>20. LIMITATION OF ABSTRACT</b> UU

NSN 7540-01-280-5500

 Standard Form 298 (Rev. 2-89)  
 Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

**Approved for public release; distribution is unlimited**

**VIRTUALIZATION OF AEGIS: A STUDY OF THE FEASIBILITY OF  
APPLYING OPEN ARCHITECTURE TECHNOLOGY TO THE SURFACE  
NAVY'S MOST COMPLEX AUTOMATED WEAPON SYSTEM**

Erik S. Roberts  
Lieutenant, United States Navy  
B.S., American Military University, 2009

Submitted in partial fulfillment of the  
requirements for the degree of

**MASTER OF SCIENCE IN INFORMATION TECHNOLOGY MANAGEMENT**

from the

**NAVAL POSTGRADUATE SCHOOL**  
**September 2011**

Author: Erik S. Roberts

Approved by: Man-Tak Shing  
Thesis Co-Advisor

Albert Barreto III  
Thesis Co-Advisor

Dan Boger  
Chair, Department of Information Sciences

THIS PAGE INTENTIONALLY LEFT BLANK

## ABSTRACT

Rising costs of proprietary equipment in legacy electronic applications are increasingly drawing resources from vital programs. Growing interest in evaluating Open Architecture technology to replace closed systems is evidenced by the number of recent publications on the subject. Researchers have approached this topic from various angles, including lifecycle management, risk simulation, total cost of ownership, and knowledge-value added measures.

This exploratory study uses open architecture hardware employing virtualization technology to test the feasibility of replacing legacy components of military systems. Virtualization has the potential to provide significant cost savings in terms of procurement, daily operation, and maintenance. Additionally, virtualization provides functional benefits such as load-balancing, greater processor utilization and storage flexibility, streamlined scalability, and simplified disaster recovery strategies.

This thesis is original research in the form of a proof-of-concept study. It details performance results of a locally-constructed test platform, designed to simulate a portion of the U.S. Navy's AEGIS Combat System. The scope of this work is to test the viability of using commodity-based hardware to achieve performance levels equal to, or greater than, current proprietary systems. Non-procurement cost comparisons are applied to the test platform and typical AEGIS systems. While this study specifically targets AEGIS, the results can be generalized to non-military legacy applications.

THIS PAGE INTENTIONALLY LEFT BLANK

## TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	MOTIVATION .....	1
B.	VIRTUALIZATION VS. CLOUD COMPUTING.....	1
C.	FOCUS OF THESIS .....	2
D.	APPROACH OF THESIS.....	3
E.	ORGANIZATION OF THESIS .....	3
II.	LITERATURE REVIEW .....	5
A.	SERVER ARCHITECTURES .....	5
B.	VIRTUAL MACHINES .....	6
C.	NAVY VIRTUALIZATION MODEL .....	7
D.	NPS RESEARCH.....	9
III.	AEGIS COMBAT SYSTEM.....	11
A.	BACKGROUND .....	11
1.	Ship Variants.....	11
2.	Combat System Variants.....	12
3.	Fleet Modernization Plan .....	14
B.	OPEN ARCHITECTURE.....	14
1.	Acquisition Considerations .....	14
2.	Transition Initiatives .....	15
3.	Technological Shortfall.....	16
C.	OPERATIONAL READINESS TEST SYSTEM .....	18
1.	Functional Description .....	18
2.	Performance Parameters.....	21
3.	Research Selection .....	22
4.	Program Software.....	22
IV.	VIRTUAL SERVER.....	25
A.	DESIGN CONCEPT.....	25
B.	SERVER CONSTRUCTION.....	26
C.	HARDWARE PERFORMANCE.....	28
1.	Physical Characteristics .....	28
2.	Scalability.....	28
3.	Redundancy .....	29
4.	Power Distribution.....	31
5.	Cooling .....	34
6.	Thermal Control .....	36
7.	Remote System Management.....	36
8.	Storage Array Network (SAN) .....	37
D.	SOFTWARE PERFORMANCE .....	37
1.	Data Traffic Management.....	37
2.	Unicast Switching.....	38

3.	<b>Jumbo Frames.....</b>	<b>38</b>
4.	<b>Virtual Local Area Networks (VLANs) .....</b>	<b>38</b>
5.	<b>Storage Management.....</b>	<b>39</b>
6.	<b>Scalability Measures .....</b>	<b>39</b>
7.	<b>Application Validity.....</b>	<b>42</b>
<b>V.</b>	<b>CONCLUSION .....</b>	<b>44</b>
A.	<b>SUMMARY OF WORK.....</b>	<b>44</b>
B.	<b>KEY FINDINGS AND CONTRIBUTIONS.....</b>	<b>45</b>
C.	<b>EVALUATION AND FEEDBACK .....</b>	<b>46</b>
D.	<b>RECOMMENDATION FOR FUTURE WORK.....</b>	<b>47</b>
<b>LIST OF REFERENCES .....</b>		<b>50</b>
<b>INITIAL DISTRIBUTION LIST .....</b>		<b>52</b>

## LIST OF FIGURES

Figure 1.	Baseline 7 (After Filz, 2009) .....	13
Figure 2.	Baseline 7, Phase 1 (After Filz, 2009) .....	13
Figure 3.	Evolution of AEGIS computer architecture.....	15
Figure 4.	Baseline 8 Logical Topology (After Filz, 2009).....	18
Figure 5.	ORTS and AEGIS Weapon System Interfaces (From Brazet, 1994) .....	20
Figure 6.	ORTS Operation (After Brazet, 1994).....	21
Figure 7.	AEGIS-VM Blade Server .....	26
Figure 8.	Blade Server Connection Diagram .....	27
Figure 9.	High Speed I/O Architecture (From Loffink, 2008).....	30
Figure 10.	Power Redundancy Modes (From Loffink, 2008).....	33
Figure 11.	Server Module Cooling Air Profile (From Loffink, 2008) .....	35
Figure 12.	I/O Module Cooling Air Profile (From Loffink, 2008) .....	35
Figure 13.	Aggregate I/O Throughput (From VMware, 2008) .....	40
Figure 14.	Average I/O Latency (From VMware, 2008) .....	40
Figure 15.	Aggregate Throughput of Multiple I/O-Intensive Virtual Machines (From VMware, 2008) .....	41
Figure 16.	Average Latency of Multiple I/O-Intensive Virtual Machines (From VMware, 2008) .....	42

THIS PAGE INTENTIONALLY LEFT BLANK

## LIST OF TABLES

Table 1.	MVME2600 series microprocessors.....	22
Table 2.	Physical Specification Comparison .....	28
Table 3.	Server Module Options (After Loffink, 2008).....	31
Table 4.	PDU Options (After Loffink, 2008).....	34

THIS PAGE INTENTIONALLY LEFT BLANK

## LIST OF ACRONYMS AND ABBREVIATIONS

ACS	AEGIS Combat System
ACTS	AEGIS Combat Training System
ADS	AEGIS Display System
ALIS	AEGIS LAN Interconnect System
AWS	AEGIS Weapon System
BMC	Baseboard Management Controller
BMD	Ballistic Missile Defense
CANES	Consolidated Afloat Networks and Enterprise Services
CEC	Cooperative Engagement Capability
CG	Cruiser, Guided Missile
CGM	Cruiser Modernization
CMC	Chassis Management Controller
COTS	Commercial Off-The-Shelf
CPU	Central Processing Unit
C2P	Command and Control Processor
C&D	Command and Decision
DDG	Destroyer, Guided Missile
DoD	Department of Defense
FCS	Fire Control System
FDDI	Fiber Distributed Data Interface
FLOPS	Floating Point Operations per Second
FMP	Fleet Modernization Plan
GB	Gigabyte

Gb	Gigabit
GbE	Gigabit Ethernet
Gb/s	Gigabits per Second
GHz	Gigahertz
GT/s	Gigatransfers per Second
HA	High Availability
HDD	Hard Disk Drive
HPCC	High Performance Computing Cluster
iDRAC	Integrated Dell Remote Access Controller
I/O	Input/Output
IOM	Input/Output Module
iKVM	Integrated KVM
iSCSI	Internet Small Computer System Interface
IWS	Integrated Warfare System
KBps	Kilobytes per second
KVA	Knowledge Value Added
KVM	Keyboard/Video/Mouse (analog switch module)
LAN	Local Area Network
LOM	LAN on Motherboard
LUN	Logical Unit (storage)
MBps	Megabytes per second
MCE	Mission Critical Enclosure
MILSPEC	Military Specification
MSEC	Milliseconds
NIST	National Institute of Standard and Technology

OA	Open Architecture
ORTS	Operational Readiness Test System
OS	Operating System
O&S	Operation and Support
PEO	Program Executive Office
PDU	Power Distribution Unit
RFI	Request For Information
RO	Real Options
SAN	Storage Area Network
SM	Standard Guided Missile
SPAWAR	Space and Naval Warfare Systems Command, Dahlgren
SPY	AN/SPY-1D Phased Array Radar System
SSH	Secure Shell
SSL	Secure Sockets Layer
TCO	Total Cost of Ownership
USMC	United States Marine Corps
USN	United States Navy
VM	Virtual Machine
vKVM	Virtual KVM
VLAN	Virtual Local Area Networks
vMedia	Virtual Media
VLS	Vertical Launch System (guided missile)
VMFS	Virtual Machine File System
WCS	Weapons Control System

THIS PAGE INTENTIONALLY LEFT BLANK

## ACKNOWLEDGMENTS

The author would like to extend his utmost gratitude to Dr. Man-Tak Shing and Mr. Albert “Buddy” Barreto for their mentorship, continued guidance, and infinite patience over the past year. Also, to a host of other faculty members such as Dr. Karl Pfeiffer, Dr. Tom Housel, and Mr. Glenn Cook, whose contributions to this writing have added immeasurable value. And to the Program Executive Office for Integrated Warfare Systems at Space and Naval Warfare Systems Command, Dahlgren for financially supporting this project.

To my beautiful wife, Trese, I truly am grateful for all of your continued support, patience, and tolerance. I love you very much and understand that personal sacrifice was not mine alone. Without you, this achievement would not be possible. It was your commitment and understanding that enabled me to devote so much of myself to this project. There were many times during the past two years that my focus was away from home. You endured all of it, and for that, I express my heartfelt gratitude.

THIS PAGE INTENTIONALLY LEFT BLANK

## I. INTRODUCTION

### A. MOTIVATION

For the past decade, forward-thinking organizations have been migrating their computer assets from legacy networks to virtualized architectures. *Virtualization* is a method by which computer resources are made more energy-efficient through real-time sharing of processor, memory, and storage assets. In a virtualized environment, operating systems and applications are no longer dependent on specific hardware configurations. In contrast to traditional computer systems where software resources are delivered through static local connections, virtualization allows for the apportioning of resources dynamically, as needed, and with finer granularity.

This thesis examines the feasibility of replacing legacy military systems with virtual machines using performance, scalability, and efficiency validity metrics. Distributed, highly scalable software resources—allocated using cloud computing technologies—facilitate a dynamic communication between a broad range of system subprocesses without requiring compromises to reliability or performance. The proposed design change outlined in this thesis applies the unique benefits of cloud computing architectures, such as dynamic information distribution and automatic scaling, to the computational-intensive AEGIS Combat System.

### B. VIRTUALIZATION VS. CLOUD COMPUTING

There has been much discussion in the press and trade papers about virtualization and cloud computing. However, they are quite different, and the two terms should not be used interchangeably. The National Institute of Standard and Technology (NIST) defines cloud computing as “a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction” (Mell & Grance, 2009, p. 1).

Virtualization, on the other hand, is an abstraction of computer resources to allow a single physical resource such as a server, an operating system, an application, or storage device appear to function as multiple logical resources; designed to deliver on-demand resources to specific users, it is but one method for delivering the on-demand resources that customers need. This is somewhat analogous to the Internet and the World Wide Web, where the Web comprises a subset of all services available to users via the Internet.

Perhaps the greatest advantage of virtualization, particularly as it pertains to shipboard environments, is the reduced physical footprint of the architecture and the optimization of existing hardware resources. As virtual machines eliminate the need for large computers that dominate interior spaces, the production of heat and noise are greatly reduced, as well as the energy requirements to power and cool them.

### **C. FOCUS OF THESIS**

This thesis focuses on developing a virtualized server rack that will match or exceed the performance of current technology employed in the Navy's AEGIS Combat System. Additionally, comparisons of non-procurement costs and values between proprietary and open source technologies will be explored. The objective of the thesis is to develop a reasonable case specific to server architecture, virtualization methodology, performance requirements, and other considerations that must be taken into account to ensure as high a probability for success as possible. This work is guided by the following questions:

- What are the technical challenges and potential benefits associated with replacing proprietary computing architectures with open source components in high-value military or commercial applications?
- What are the software and hardware requirements to best address those challenges and/or realize those benefits?
- What measures of efficiency, performance, etc. may be used to practically assess the value of using open source technology versus traditional closed source systems?

#### **D. APPROACH OF THESIS**

A working relationship was developed between the Virtualization Lab at Naval Postgraduate School (NPS) and the Program Executive Office for Integrated Warfare Systems (PEO IWS 7) at Space and Naval Warfare Systems Command (SPAWAR), Dahlgren. Through this relationship, research was conducted on existing commercial off-the-shelf (COTS) solutions and implementations, including virtual servers. There was also a thorough review of network applications, Web applications, and Internet architectures to determine which applications and architectures might best suit a shipboard military environment.

Based on input from SPAWAR, Dahlgren, the AEGIS system was categorized into several subsystems such as command and control, sensor interface, weapon interface, operator diagnostics, and training. Due to time constraints, only the diagnostic subsystem was targeted as the focus of this project. The AEGIS diagnostic module is referred to as ORTS (Operational Readiness Test System) and does not demand the same security considerations as other functional units of the AEGIS system. From this premise, a testing platform comprised solely of open source hardware and software components was designed and constructed. This virtual server became the vehicle from which comparison performance data was collected for the purposes of this project. It is believed that the general-purpose applicability of this approach makes it ideal to serve as the basis for future work in this, as well as other, fields of study.

#### **E. ORGANIZATION OF THESIS**

The remainder of this thesis is organized as follows:

Chapter II details the background and research for this thesis. Open Architecture technologies are examined as a viable solution to scaling or replacing legacy closed system architectures with the potential to address efficiency problems inherent to the proprietary domain. Virtualization is examined in depth as a potential platform on which to develop such solutions. The CANES program, a current Navy IT initiative to improve shipboard IT environments through virtualization, is explored.

Chapter III reviews the evolutionary development of past and current AEGIS baselines and their applicable computing architectures. Functional and non-functional values of legacy components are evaluated in terms of utility and efficiency based on information and research captured in Chapter II.

Chapter IV details performance measures collected from the testing platform and draws comparisons of performance and other value metrics with typical AEGIS shipboard systems based on specifications detailed in Chapter III. VMware® virtual software, running on Intel-based Dell® blade servers is proposed as a technical solution.

Chapter V summarizes the thesis, identifies its key contributions, and outlines a framework for future work.

## II. LITERATURE REVIEW

### A. SERVER ARCHITECTURES

Over the years, accepted server architecture has oscillated between centralized and decentralized models. In this context, reference is made specifically to *server* architecture and not to the more general *system* architecture, in which a broad scope is applied to interfaces among all components or subsystems, and the interface between the system as a whole and the external environment. The narrower scope of server architecture has been selected to best meet the project goals targeted within this thesis.

The concept of centralized server architecture began with mainframe computers. Many businesses required the execution of heavy calculations to maintain their position with emerging trends and possibly to gain a competitive advantage. Large systems, called mainframes, were introduced that could accomplish calculations that were far beyond the scope of what was possible with a single terminal system. By hardwiring relatively simple user terminals to the mainframe's computers, significant computing power could be distributed to the system terminals throughout an organization.

These early “virtual” machines, such as the IBM 370, allowed excess resources, primarily measured in CPU cycles, to be used more efficiently (Note: floating point operations per second or “FLOPS” are generally considered a more accurate measure of comparing improvements in processor utilization over time). Large corporations were now able to run dozens, even thousands, of terminal machines from a single mainframe computer. However, at this time, most enterprises did not have the need for a multi-million dollar system that could run over two hundred times the number of operations required to meet business output. These companies instead allocated their budget and design resources to less capable yet less expensive commodity servers.

As the number of transistors able to be placed on a circuit board doubled every eighteen months, as represented by “Moore's Law” (Moore, 1965), mainframe technology gave way to popular commodity server components that could satisfactorily

accomplish most business goals using less expensive hardware. Although the amount of data that required increasingly more effective processing grew over time, the expansion of business output requirements was outpaced by the speed of emerging individual computer technology. The significant improvements in the computational speed of these “stand-alone” systems allowed almost any company to meet their market demand through the use of *decentralized* commodity server architecture.

This form of decentralized computing with less expensive commodity servers allowed a firm to spend far less money in the acquisition and maintenance of its server architectures. Firms were able to purchase hundreds of commodity servers for less than half the cost of a single mainframe computer. Soon, the ultimate computing power provided by mainframes was reserved for only the heaviest of computational workloads or when multiple operating systems were required to work in parallel.

However, as technology continued to advance, computing resources soon outpaced the technical needs of the business. The better, faster, cheaper aspects of inter-networking amongst distributed locations allowed server resources to be pooled and accessed remotely. Servers were then designed to exploit the abundance of unique decentralized systems distributed across an enterprise. This organizational bloat led to increased expenditures for labor, travel, and maintenance to manage and support system hardware. Local processing was no longer such an inexpensive solution for enterprises.

The economic curve of technology progression has come full circle; centralized computing is once again the more economical choice for most organizations. In short, the speed of the computer has outpaced that which is required by business—and improvements in the speed of packet transfers across networks has further reduced the need for decentralized systems.

## **B. VIRTUAL MACHINES**

In contrast to the hardwired “virtual” machines of the mainframe era, modern virtual machines (VM) use software applications to accomplish the interconnectivity required for resource sharing between servers, or *hosts*, and their terminals, called *clients* or *guests*. However, while virtualization software allows a host computer to create and

run multiple virtual environments—which are then accessed through client computers—it is important to note that the type of virtualization can take on several forms. For example, *storage virtualization* refers to the process of abstracting logical storage capabilities from an actual physical storage device. At the other end of the spectrum, virtualization software may be used to emulate a server’s entire computer system. This is helpful if the client’s operating system is different than that of the server. For example, virtualization allows a Linux OS that natively resides on a server to be run as a guest on top of whatever resident OS the client happens to be running, such as Microsoft Windows.

*Server virtualization* is accomplished through a software application that allows for the logical division of a physical server—usually enhanced with multiple state-of-the-art processors, expandable storage capacity, and large amounts of RAM—into multiple, unique virtual environments. These virtual environments reserve space on the system in logical files that emulate physical pieces of hardware, running whatever OS or application it is configured to run. For the purposes of this research, we chose VMware® products to provide the application software to host our virtual server. This choice was made primarily due to the fact that VMware® had an established contract with the U.S. Marine Corps for deployment and support of server and desktop applications.

### **C. NAVY VIRTUALIZATION MODEL**

Although there were no existing examples to model our virtual AEGIS server after, various components within the military have made great strides with other applications of virtualization technology. In fact, all service branches have begun to employ virtualized environments for many fixed IT architectures, principally with desktop computer networks. The Navy, in particular, has sought to craft an IT solution for mobile environments. The Consolidated Afloat Networks and Enterprise Services (CANES) program is the Navy’s latest effort to improve shipboard networks by consolidating five legacy systems currently in use aboard Navy ships.

CANES is designed to create a shared resource base for several functional areas or *domains*, including command & control, communications, computers, intelligence, surveillance, and reconnaissance (collectively known as C4ISR). The CANES

virtualization solution replaces disparate computer networks that encompass unique architecture specifications and distinct incompatible resources with a single hardware infrastructure to replicate and disseminate (i.e., *virtualize*) a suite of common-access software applications. A recent article in Defense Systems magazine quoted Captain D.J. LeGoff, program manager for the Tactical Networks Program Office, as saying, “...the CANES program...validates technology maturity...the program foundation is built upon cost containment, open architecture, and competition throughout the program’s lifecycle” (Corrin, 2011).

Captain Kevin Hooley, assistant chief of staff for readiness and training at Navy Cyber Forces, also advocates the merits of virtualization programs like CANES, describing how the Navy is adopting a strategy similar to the knowledge-centric approach that many Fortune 500 companies have embraced in recent years. “This is very exciting. We’re finally understanding the value and speed of information in the 21st century” (Richfield, 2010).

The Navy and industry recognize that because knowledge is stored securely on the IP network instead of in a server or hard drive, computers can be purchased on the open market in quantity and from multiple suppliers. And, as hardware and software advance, components can be quickly upgraded without taking the system offline. Interoperability remains the greatest problem, Hooley said, “The various applications and virtual systems must be coherent so they don’t defeat each other. This can only get worse as the amount of hardware and software for traditional C4ISR functions and now combat systems grows” (Richfield, 2010).

The ultimate value of this approach lies in the fact that it does not require reinventing the wheel; nothing that currently exists within a particular legacy system needs to be re-developed. Mike Twyman, vice president of integrated command, control, communications, and intelligence systems at Northrop Grumman’s Information Systems unit says that, “With [this] approach, the Navy doesn’t have to rip out a rack to install a new processor—they just add to it. Leveraging COTS technology against the custom technology that’s been used in the past is a much stronger position that reduces the total ownership cost” (Richfield, 2010).

#### **D. NPS RESEARCH**

At the Naval Postgraduate School, faculty and students are exploring cost-effective ways of employing virtualization and cloud technology. In applications where exceptional performance and high availability are not primary goals, first-generation servers and storage devices are being reutilized. Several brands of virtual server operating systems (UNIX®, Microsoft®, Solaris®, etc) running on Intel® hardware have been used to help make informed decisions regarding what technologies might best serve various missions. To that end, thesis students and their advisors are researching ways in which virtualized systems might be used to reduce complexity and load on the warfighter, and improve connectivity with remote devices that require access to servers over bandwidth-constrained links.

Previous thesis research, specific to AEGIS, has investigated the applicability of using open architecture (OA) methodology in support of anticipated future maintenance requirements and functional upgrades. Captain Joseph Uchytil, USMC, first approached the subject of improving the way the Department of Defense (DoD) made decisions concerning the integration of current and future processes and systems through the application of Knowledge Value Added (KVA) methodology. His June 2006 thesis focused on the processes involved in track management aboard AEGIS platforms.

In June 2007, Ensign Jameson Adler, USN, and Ensign Jennifer Ahart, USN, produced a thesis that used KVA to estimate performance improvements by employing an OA approach to AEGIS software upgrades. The basis of their research was to measure how an OA approach might affect the software upgrade and maintenance process for the AEGIS Integrated Warfare System (IWS) since the majority of total lifecycle costs in IWS acquisitions occur during the Operation and Support (O&S) phase.

This research was expanded in October 2007 by NPS Professors Tom Housel and Johnathan Mun through a proof-of-concept case study that quantified the benefits of OA within the AEGIS software maintenance and upgrade processes through a multi-phased framework of Knowledge Value Added and Real Options (KVA+RO) in order to provide

decision-makers with a systematic approach for analyzing benefits and assessing risks of potential technological acquisitions.

Navy Lieutenant Sylvester Thompson's thesis, March 2008, endorsed the need to adopt an OA approach to guide the replacement of aging AEGIS components to offset the costs of time-consuming maintenance and upkeep. Thompson's thesis transitioned from an overarching perspective of the benefits of OA to an investigation into the performance parameters of COTS alternatives for a specific proprietary component of the current AEGIS system.

Finally, Captain Luis Tiglao, USMC, researched the application of VM technology to create models and/or simulations (M&S) of current IT capabilities used by military operating forces. His thesis focused on the cost savings and operational benefits of virtual environments in the application of M&S and applied it to the DoD Verification, Validation, and Accreditation (VV&A) process.

While the research papers outlined in this section advanced understanding of the technology and provided a springboard for this project, they did not consider how virtualization technology might be applied to subsystems within the AEGIS Combat System. Therefore, an original approach was required to approximate the computing capability of AEGIS using OA hardware and software components.

### III. AEGIS COMBAT SYSTEM

#### A. BACKGROUND

##### 1. Ship Variants

Initially contracted in 1969, the first Navy warship equipped with the AEGIS Weapon System (AWS) was USS *Ticonderoga* (CG 47). Commissioned in January 1983, and subsequently decommissioned in September 2004, *Ticonderoga* led the development, acquisition, and construction of 27 *Ticonderoga*-class cruisers (CGs) and 61 *Arleigh Burke*-class destroyers (DDGs). As of today, there are 83 of 88 AWS-equipped U.S. Navy ships in service. However, current planning calls for the production of additional DDGs. These AEGIS ships were built in four fundamental *baselines*:

- Baseline 1 CG 47 through CG 51 (all decommissioned)
- Baseline 2 CG 52 through CG 58
- Baseline 3 CG 59 through CG 64
- Baseline 4 CG 65 through CG 73 and all DDGs
- DDG Flight I DDG 51 through DDG 71
- DDG Flight II DDG 72 through DDG 78
- DDG Flight IIA DDG 79 and following

(Source: AEGIS Combat System)

By building DDGs in *flights*, incremental development allowed for technological advances to be incorporated during construction. Flight II, introduced in FY92, incorporated improvements to the AN/SPY-1 phased-array radar and the SM-1 standard missile and launcher. Flight IIA was introduced in FY94 and added a helicopter hangar.

All AEGIS baselines and flights are functionally similar with differences related to specific equipment such as military specification (MILSPEC) computers and peripherals, processing techniques such as COTS processors, LAN interface protocols, physical point-to-point interfaces, and individual architectural models as characterized by proprietary executive program design and operating environments.

## 2. Combat System Variants

Building on the initial four baselines that defined AEGIS ship construction, upgrades to the AEGIS Combat System (ACS) programming were also defined in terms of baselines. There are two active programming baselines:

- AEGIS 3A/5.3.x Standard AWS baseline
- BMD 3.6.x Ballistic Missile Defense
- BMD 4.0.1 *future release*

(Source: AEGIS Combat System)

Specific to AEGIS computing capability, Baseline 5 served to upgrade the Command and Control Processor (C2P) and various data link information systems. Baseline 6, Phase I incorporated COTS hardware into the Fiber Distributed Data Interface (FDDI) Local Area Network (LAN) and into the adjunct computer for the AEGIS Display System (ADS). Baseline 6, Phase III introduced Ballistic Missile Defense (BMD) and Cooperative Engagement Capability (CEC) without COTS components in either program. Baseline 7, Phase 1C was the first AEGIS baseline to finally implement an actual open architecture strategy in 2005.

The AEGIS Weapon System consists of ten basic subsystems, including:

- AEGIS Combat System (ACS)
- Command and Decision System (C&D)
- Phased Array Radar System (SPY)
- Weapons Control System (WCS)
- Fire Control System (FCS)
- Guided Missile Vertical Launching System (VLS)
- Standard Guided Missile (SM)
- AEGIS Display System (ADS)
- Operational Readiness Test System (ORTS)
- AEGIS Combat Training System (ACTS)

Technically, the ACS is a subsystem of the AWS, including the computing architecture and excluding weapons systems and interfaces. However, the terms are often used interchangeably. Figure 1 and Figure 2 illustrate the fundamental differences in signal flows between a static and dynamic AEGIS LAN Interconnect System (ALIS).

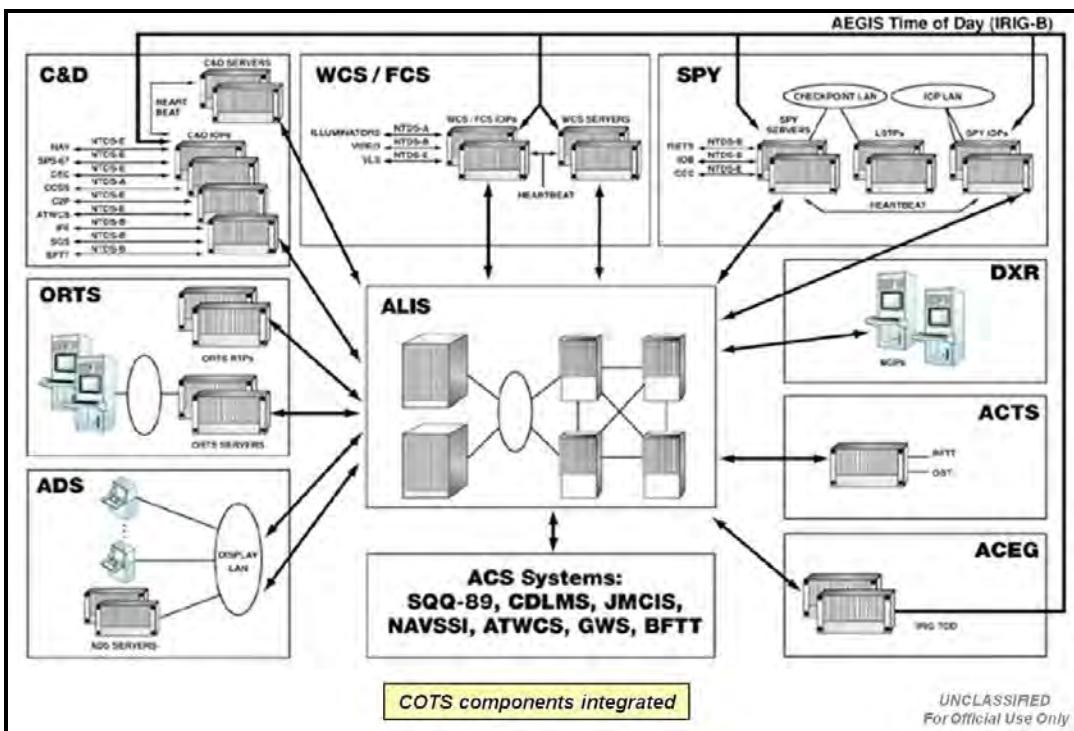


Figure 1. Baseline 7 (After Filz, 2009)

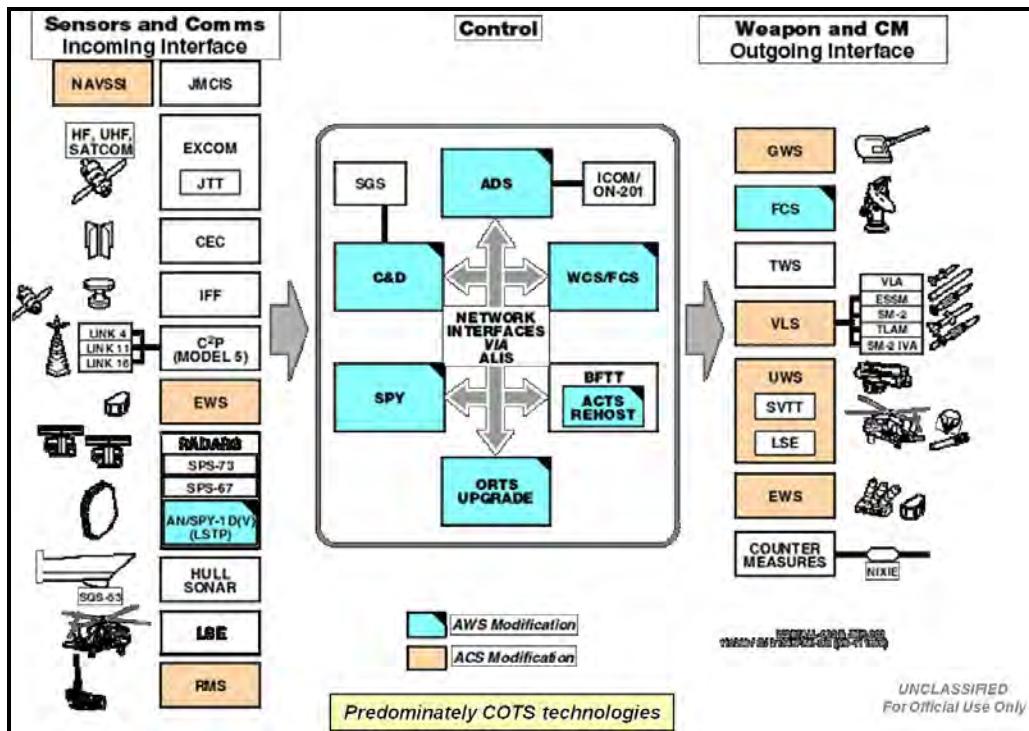


Figure 2. Baseline 7, Phase 1 (After Filz, 2009)

### **3. Fleet Modernization Plan**

Modernization of the combat system architecture for existing cruisers and destroyers remains an ongoing process, with the goal to create a more modern networked computing environment. Unfortunately, while upgrades implemented through the AEGIS Fleet Modernization Plan (FMP) increasingly employ additional commercial strategies, the focus appears to be mainly on integrating COTS hardware components into static legacy architectures. While utilization of componentized software enables some reuse between ships, it remains proprietary technology, and does little to uncouple embedded software applications from specific hardware subsystems.

Moving forward, beyond retrofitting older ships, the Navy has decided to restart production of the *Arleigh Burke*-class destroyers. The future DDGs will incorporate new computing technologies such as the SPY-1D(V) Multi-Mission Signal Processor (MMSP) and the BMD 4.0 programming baseline, as well as improved detection and processing capabilities in electronic warfare (EW) and undersea warfare (USW) suites.

## **B. OPEN ARCHITECTURE**

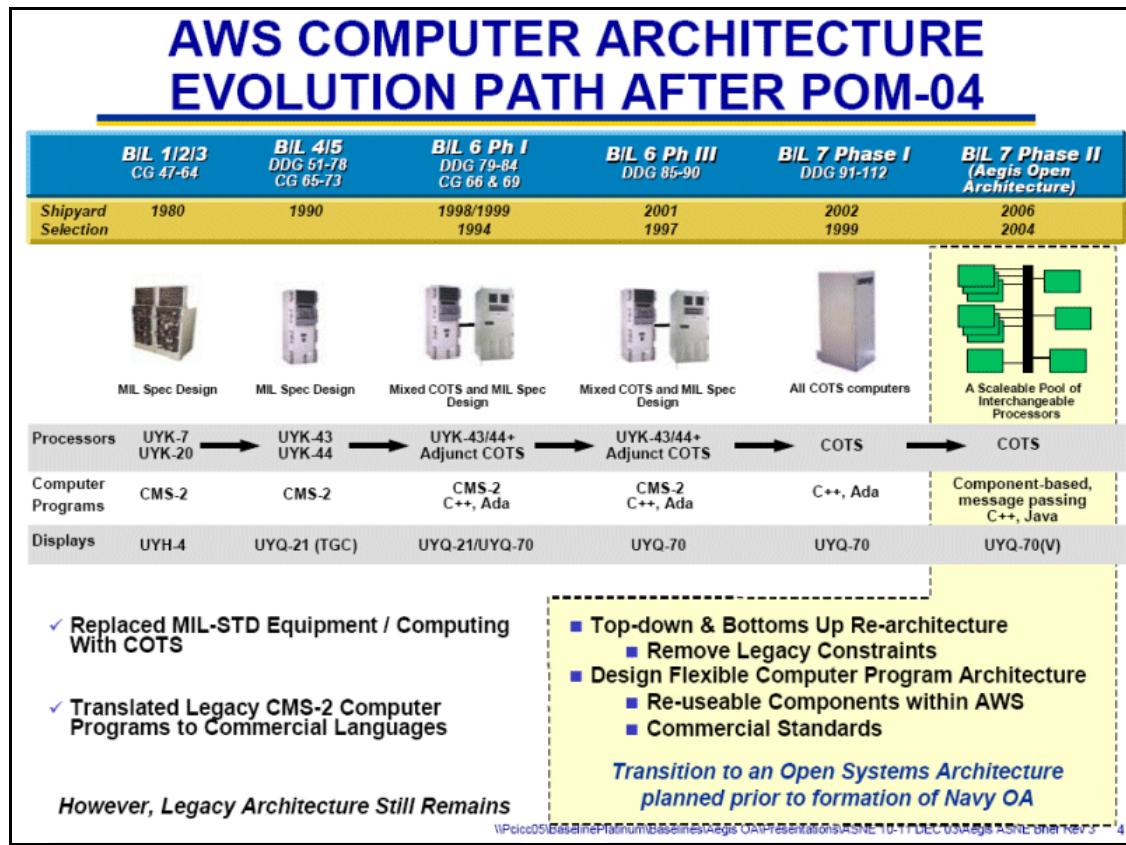
### **1. Acquisition Considerations**

The U.S. Navy has historically acquired systems that are proprietary in nature and which require sole source components and maintenance services to support them. This practice employs a very limited number of suppliers and, in turn, causes procurement, repair, and modernization initiatives to become unnecessarily expensive. Additionally, the development and implementation timelines of proprietary/sole source military systems can take 15 years or more before a needed capability reaches the warfighter.

Fortunately, the Navy in general has been on a transition path to open architecture since 2002, transforming its traditional business practices toward an open product-line strategy that will continue to distance service requirements from proprietary technology and contractors, and to better leverage the available network of developers for delivery of rapid, cost-effective acquisition solutions.

## 2. Transition Initiatives

The AWS is responsible for allowing Navy warships to detect, track, and prosecute multi-mission contacts and targets. The computer-based command and decision (C&D) element is at the core of the ACS. This interface is absolutely critical to providing AEGIS the ability to execute simultaneous resource-intensive operations in order to be effective against real-time airborne, surface, and subsurface threats. Figure 3 illustrates the evolution path of the AEGIS computer architecture, culminating with the Navy's FY04 plan to transition to a component-based pool of interchangeable processors.



Source: AEGIS ASNE Brief (Williams, 2003, p.4)

Figure 3. Evolution of AEGIS computer architecture

AEGIS continues to transition toward a genuine open architecture strategy. However, it is a slow progression. Despite steps taken to leverage OA strategies to further innovate the AEGIS computing infrastructure within the 2008 Cruiser

Modernization (CGM) Baseline, the Navy concurrently proposed to the Senate Armed Services Committee a five-year deal to continue AEGIS development via sole source contracting.

The request reflects the challenges associated with moving large and complex legacy systems toward open competition due to long-standing proprietary development. The committee approved the Navy's request only under the provision that a plan be created to provide for—and verify—steady, incremental progress toward opening AEGIS.

The committee directed that no greater than 50 percent of the amounts authorized for fiscal year 2009 may be obligated under sole source contracts, until 30 days after submission by the Secretary of the Navy of a detailed program plan for implementing OA for the AEGIS combat system. The program plan shall be included in subsequent quarterly reports to the congressional defense committees on Naval Open Architecture, and shall include methodology and scheduling for incrementally opening the AEGIS combat system. The plan must provide for measuring discrete progress toward achieving a full open system commensurate with introduction of the 2012 AEGIS baseline (formerly referred to as 'COTS Refresh 3') ("AEGIS Open Architecture," para.8).

The total cost for the AEGIS Weapon System is \$42.7 billion with the predominant driver of cost being operations and support (O&S) at \$22.2 billion. Clearly there are opportunities for the production community to make some impact on their O&S costs even if they appear trivial. For example, a one percent reduction in [total cost of ownership] today would buy a brand new ship in 10 years. Some areas that can be targeted by the production community include repair parts, spares, and in-service engineering reductions ("AEGIS Combat System," para.14).

### **3. Technological Shortfall**

In September 2009, a Fleet Review Panel was convened by Admiral John C. Harvey, commander of Fleet Forces Command, to conduct an outside assessment into the readiness of the surface force. The seven-member panel, chaired by retired Vice Admiral Phillip Balisle, produced a report that communicated an absolute disapproval of Navy procurement decisions spanning more than a decade.

The following is an excerpt from the report:

Ships are not ordering replacement voltage regulators, which SPY radars need to help manage their prodigious consumption of ship's power. Crews aren't ordering them because technicians can't get the money to buy spares, so commanders are knowingly taking a risk in operating their AEGIS systems without replacements.

The technicians can't get the money to buy spare parts. They haven't been trained to the requirement. They can't go to their supervisor because, in the case of the DDGs, they likely *are* the supervisor. They can't repair the radar through no fault of their own, but over time, the non-responsiveness of the Navy system, the acceptance of [equipment] degradation by the Navy system and their seniors...will breed (if not already) a culture that tolerates poor system performance. The fact that requests for technical assistance are up Navy-wide suggests there is a diminished self-sufficiency in the surface force. (Ewing, 2010)

This report, while not directly admonishing the Navy's over-reliance on proprietary equipment, speaks to the same downward spiral described in this thesis. When maintainers are not permitted to receive the proper tools, including formal training, of the shipboard equipment for which they are responsible due to restrictions imposed by concerns over proprietary technology, costs to the user rise. Contractors with guarded instruments and knowledge must be flown out to meet the ships at government's expense. Often the "fix" is a matter of plugging in a special stand-alone laptop running exclusive software to reconfigure or reboot the system.

There is hope however. Figure 4 illustrates the new logical topology designed for Baseline 8 that is being implemented in conjunction with Advanced Capability Build 2008 (ACB08). Operational testing of CG 52 through CG 58 (upgraded with ACB08) began in July 2010 and is expected to be completed this year.

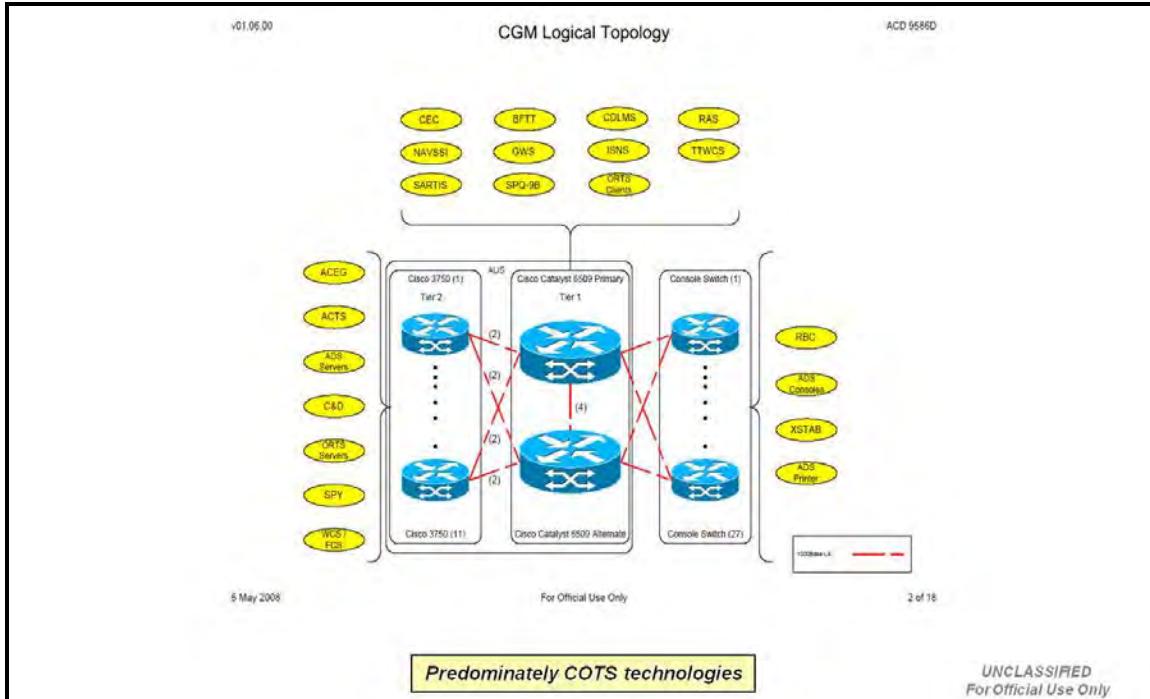


Figure 4. Baseline 8 Logical Topology (After Filz, 2009)

While this topology may not at first appear to be ground-breaking innovation, particularly from the perspective of technological advances found in the private sector, it represents a major departure from the way the Navy has been building AEGIS computing architectures.

Baselines 5.3.x and 3.6.x are predominately comprised of MILSPEC computers, peripherals, and point-to-point interfaces that have been marginally integrated with COTS technologies. Important to note is the fact that some intrusion points exist with these types of hybrid architectures. In contrast, Baselines 6.3.x and following (Flight IIA DDGs) were designed as predominately COTS constructs.

## C. OPERATIONAL READINESS TEST SYSTEM

### 1. Functional Description

The AEGIS Operational Readiness Test System (ORTS) testing module serves as a structured built-in test subsystem of the ACS to streamline shipboard maintenance and

test operations and functions as an integral part of AEGIS operations, providing fail-safe interfaces to tactical elements without degrading subsystem or system performance.

ORTS was developed concurrently with AEGIS in the 1970s to support online automated fault isolation. Design specifications included significant lifecycle cost (LCC) and Operational Availability ( $A_o$ ) analyses to ensure near-optimum allocation of diagnostic functions. Through this analysis process, ORTS became the predecessor to modern integrated diagnostic systems, providing an optimal mix of diagnostic modeling and analysis within the AEGIS test and evaluation subsystem.

Specific requirements allocated to ORTS include: comprehensive online testing, evaluation and reporting of AEGIS element status and to command and decision (C&D) echelons, central control of AEGIS system initialization, casualty reconfiguration of multiple integrated computer groups, direction and coordination of AEGIS diagnostic and maintenance activities, and procedural and functional integration with the AEGIS logistics support system (Brazet, 1994).

The original configuration for ORTS aboard early cruiser baselines is shown in Figure 5. Operations are controlled by the ORTS computer program resident in the AN/UYK-20 computer. System interfaces to AEGIS element test functions are through inter-computer channels to the three AEGIS tactical computer suites (Brazet, 1994).

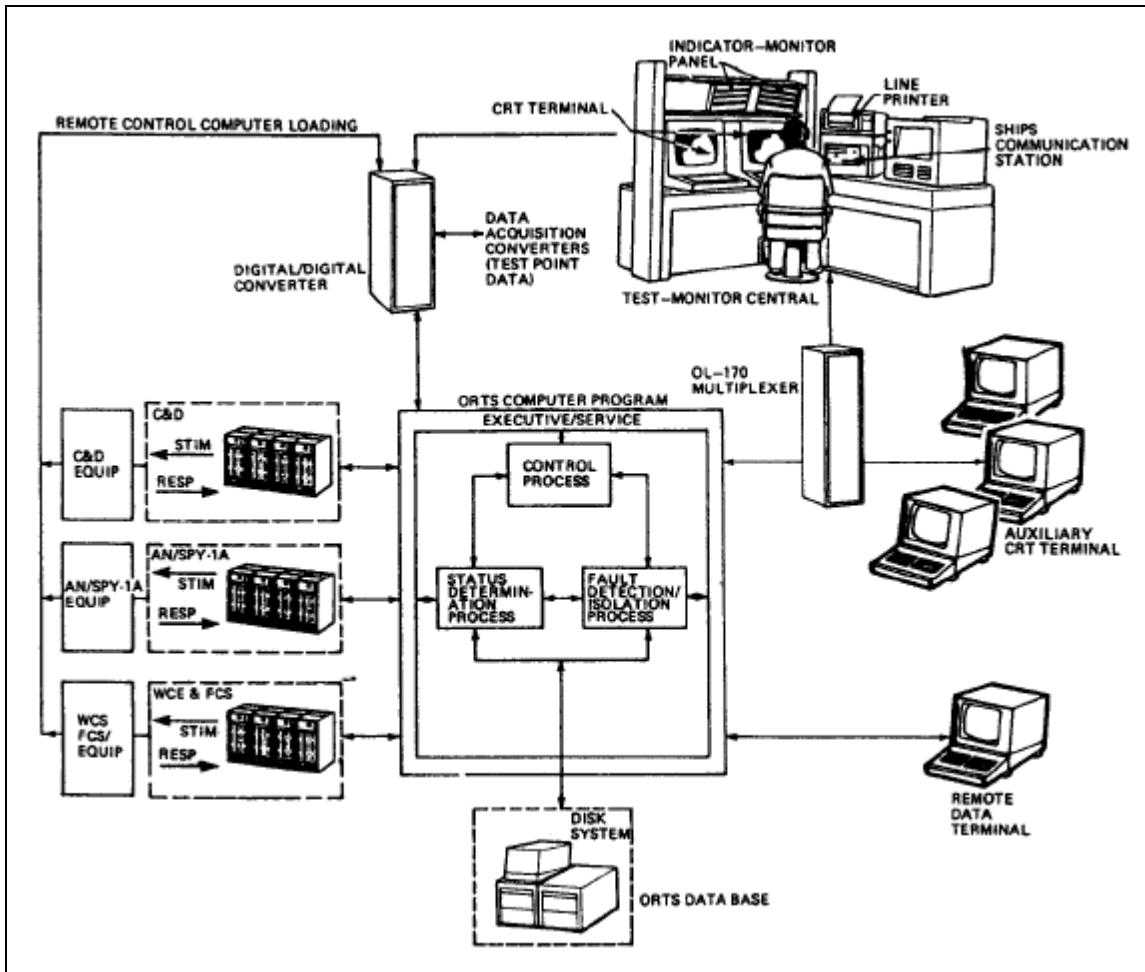


Figure 5. ORTS and AEGIS Weapon System Interfaces (From Brazet, 1994)

ORTS operations as shown in Figure 6 depict an automatic test system in which the maintenance supervisor has ultimate control of the test system. The maintenance supervisor is able to direct the scheduling and accomplishment of test operations, monitor equipment status, and direct repair efforts, as required, based upon fault information received at the console.

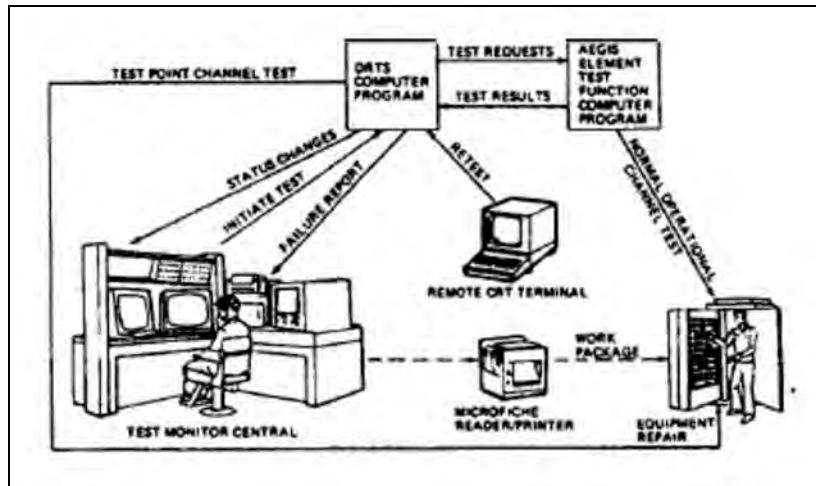


Figure 6. ORTS Operation (After Brazet, 1994)

Test results are evaluated by ORTS to determine if equipment status has changed. ORTS will then automatically display those faults with highest level of system impact. A change in equipment status, as a result of a fault condition, will prompt the initiation of fault isolation diagnostics to determine the cause. A re-test of the faulty function, once corrected, evaluates the repair and updates system status. This provides a cost effective capability to detect and isolate faults with the ultimate goal of maximizing equipment availability and minimizing total user cost.

## 2. Performance Parameters

ORTS is capable of constantly monitoring thousands of critical operating points, providing maintenance data for rapid correction or repair by the ship's crew. The ORTS architecture incorporates Motorola® 2604 PowerPC™ processors. Originally built for personal computers, PowerPC CPUs have since become popular as embedded, high-performance processors.

The MVME2600 series is a family of VME processor modules based on the Motorola® PowerPlus™ VME architecture with PowerPC™ architecture-compatible microprocessors. Two basic processor models offer either 333 MHz @ 256MB RAM or 400 MHz @ 512MB RAM. Table 1 lists the performance specifications of the Motorola VME 2600 series processor modules.

Specification	Motorola 2600 series
<b>Microprocessor Class</b>	MPC60x
<b>ECC DRAM</b>	Up to 512MB
<b>L1 cache size</b>	16KB or 32KB
<b>L2 cache size</b>	256KB
<b>Flash Memory</b>	8MB on-board, 1MB socketed
<b>PCI Mezzanine Connector</b>	64-bit
<b>PMC Expansion Slot</b>	79°F/26°C avg internal temp
<b>Serial Ports</b>	2 or 3 asynchronous 1 or 2 synchronous/asynchronous
<b>Ethernet Interface</b>	32-bit PCI local bus DMA
<b>Fast SCSI-2 Bus Interface</b>	8-bit or 16-bit

Source: [www.motorola.com/computer/literature](http://www.motorola.com/computer/literature)

Table 1. MVME2600 series microprocessors

### 3. Research Selection

ORTS was specifically targeted for this thesis to serve as a springboard for comparing performance metrics between the legacy architecture of typical AEGIS configurations and the OA architecture of a Dell® blade server. ORTS was chosen because it does not incur the same security considerations as other functional units.

### 4. Program Software

A copy of the latest version of the ORTS program (BL7.1.R) was obtained from SPAWAR in order to evaluate software performance in a meaningful way through

measures of efficiency under various conditions of resource stress. The government package was comprised of the ORTS application embedded with the Solaris 8 operating system (OS).

Solaris™ is UNIX-based OS originally developed by Sun Microsystems®, but has been operated by Oracle Corporation® since Oracle acquired Sun in January 2010. Although the Solaris OS is generally considered highly scalable, its performance is best demonstrated on SPARC™ processing systems. SPARC, short for scalable processor architecture, is a reduced instruction set computing (RISC) instruction set architecture (ISA) that was also developed by Sun.

Solaris was originally developed as proprietary software, but Sun released most of the source code in June 2005 and founded the OpenSolaris open source project. However, Oracle later decided to discontinue the OpenSolaris distribution. Solaris 8 was released in February 2000 and its product support is scheduled to end in March 2012 (Sun Microsystems, 2009).

Since the ORTS software was designed to only run on UNIX-based systems (Sun Blade 150 or later), we were unable to use it on our Intel-based Dell system. It is ironic that the project underlying this thesis, developed in part to demonstrate the utility of open source products, should be derailed by a computer program that is only compatible with older OS versions designed to run on specific types of processors.

THIS PAGE INTENTIONALLY LEFT BLANK

## IV. VIRTUAL SERVER

### A. DESIGN CONCEPT

Lockheed Martin began installing its AN/UYQ-70 Mission Critical Enclosure (MCE) aboard U.S. Navy ships approximately 10 years ago to better meet the unique housing requirements of AEGIS computers. The MCE design made good use of COTS technology that incorporated innovative subsystems to handle vibration isolation, thermal management, and regulation of input power. However, the MCE only serves to better protect equipment from the rigorous environment of AEGIS warships. With the MCE acquisition, the Navy chose not to address the fundamental problem of continuing to use antiquated computing technology.

Several measures of performance motivated the various decisions made during the design of our AEGIS-VM server, such as providing equal or greater computational performance and storage capacity, ready scalability, smaller physical footprint, reduced energy consumption and environmental requirements, and universal application across all AEGIS platforms.

A primary component to determining the applicability of using virtualization to replace any legacy application is the selection of appropriate hardware. However, there were few examples, relevant to this research, from which to draw lessons or emulate architectural design. Regardless, a physical testing platform was necessary to evaluate performance in a meaningful way through measures of efficiency that could be compared to those of typical shipboard systems. To that end, a top-of-the-line virtual server rack was created using Dell® open source products that are based on Intel® central processing units (CPUs).

Due to its much smaller size, the COTS server rack reduces the environmental footprint (i.e., physical space, power consumption, heat dissipation) while meeting or exceeding the computational capacity of the AEGIS components housed in the MCE.

## B. SERVER CONSTRUCTION

Building upon a basic Dell® PowerEdge™ M1000e blade enclosure, standard PowerEdge™ M610 server blades were upgraded with two Intel® Xeon® quad-core processors (E5540 @ 2.53 GHz) and permanent memory was supplied by a Dell® EqualLogic™ PS6000XV iSCSI Storage Array. This provided a rated bus speed of 5.86 gigatransfers per second (GT/s) and 24 gigabytes (GB) of system memory. Dell® PowerConnect™ M6220 Gigabit Ethernet (GbE) switches were “stacked” to create a single logical switch that provides seamless fault tolerance through cross-connected modules and an aggregated throughput of 4 GbE between the server rack and the IP network. Finally, server management was supplied through interfaces with an integrated Dell® Chassis Management Controller (CMC) as well as an external management console. Figure 7 is the front view of the completed AEGIS-VM testing platform and Figure 8 illustrates the interconnecting cables between the various system components.



Figure 7. AEGIS-VM Blade Server

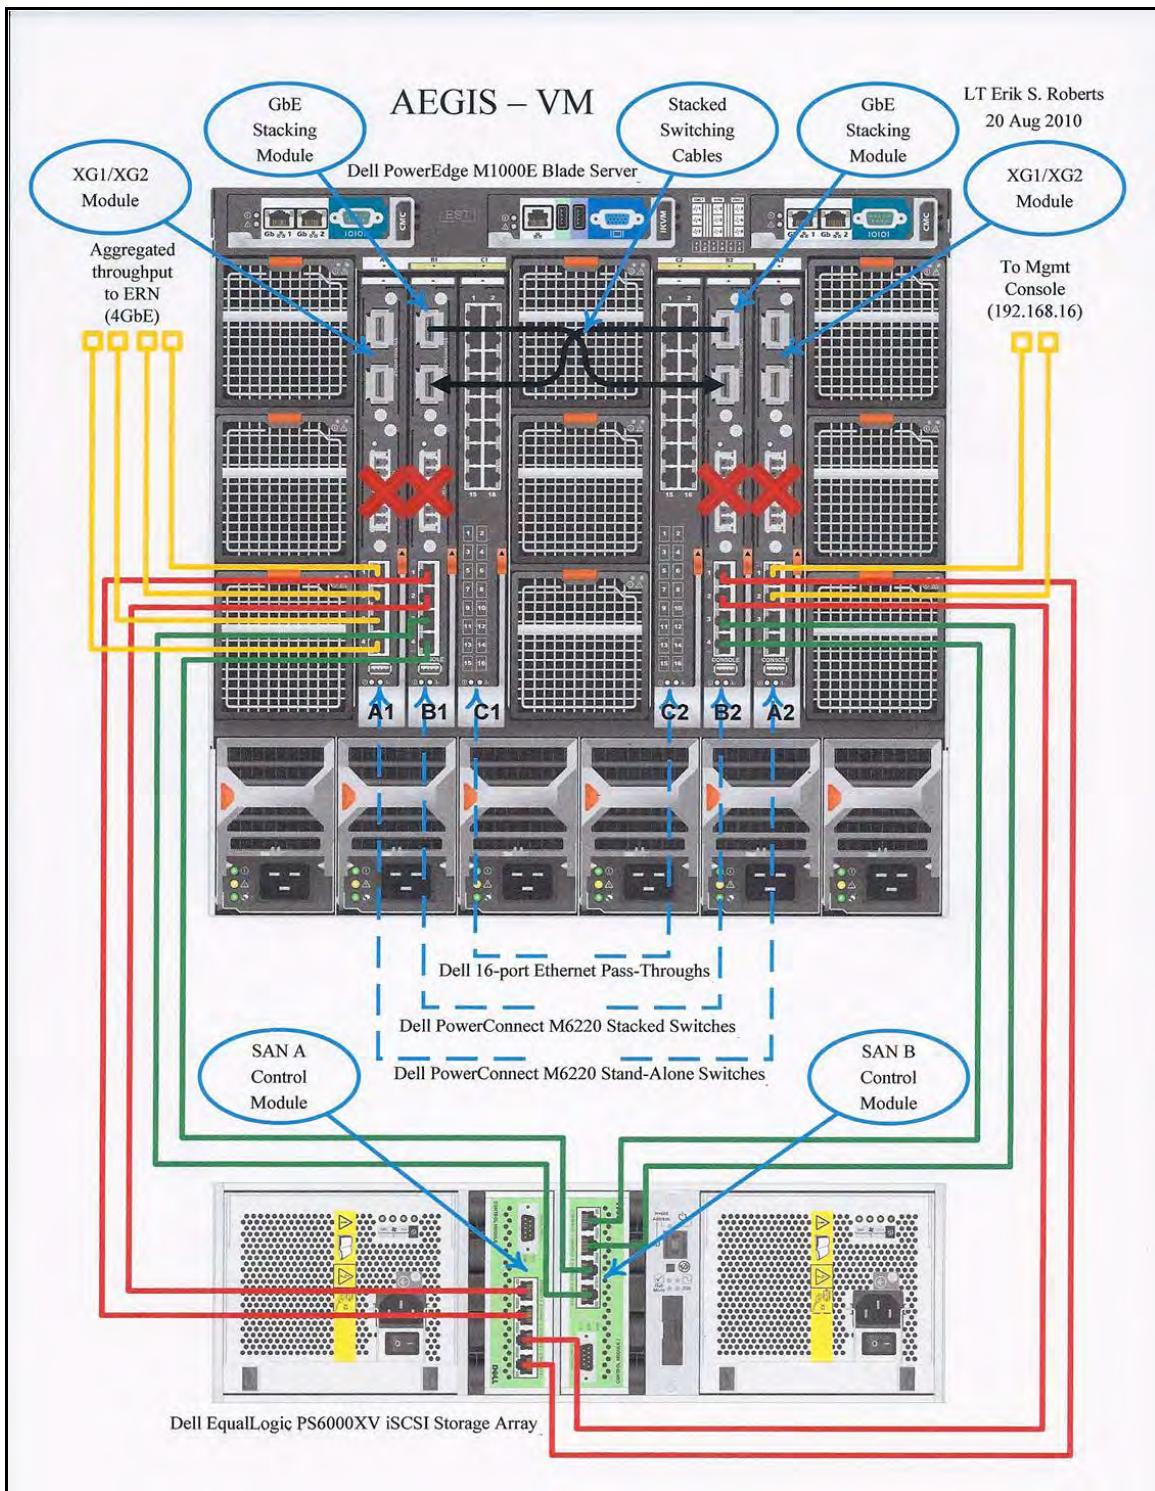


Figure 8. Blade Server Connection Diagram

## C. HARDWARE PERFORMANCE

### 1. Physical Characteristics

Table 2 details a side-by-side comparison between the physical specifications of the AN/UYQ-70 Mission Critical Enclosure and the AEGIS-VM virtual server rack.

Specification	MCE	VM
Size	28" W x 75" H x 36" D	24" W x 48" H x 36" D
Weight	2,000 lbs/907 kg	500 lbs/227 kg (appox. wt of fully populated server rack)
Thermal Regulation	Water-cooled (70°F/21°C to 76°F/24°C)	Air-cooled (70°F/21°C, avg room temp)
Internal Temperature	95°F/35°C max. internal temp	79°F/26°C avg internal temp

Source: Lockheed Martin Marketing Products: [www.Q70.com](http://www.Q70.com)

Table 2. Physical Specification Comparison

### 2. Scalability

The PowerEdge™ M1000e blade enclosure offers scalability up to the following performance values:

- Capacity for up to (16) half-height blade server modules
- Capacity for up to 64 GB of memory per blade server module
- Capacity for up to (6) network & storage input/output (I/O) interconnect modules
- Comprehensive I/O options support dual links of 20 GB/s (high-speed connectivity to storage array and the IP network)

### 3. Redundancy

Integrated Dell Remote Access Controllers (iDRAC) reside on each blade server module and are connected via fully redundant 100 Mbps Ethernet connections to dedicated 24 port Ethernet switches on two redundant (1+1) CMCs. The two Ethernet connections provide redundancy for internal system management interfaces, while exposure to systems on the IP network is channeled through the CMC's external management Ethernet interface.

Each iDRAC supports three redundant multi-lane fabrics [Dell's term for its method of encoding, transporting, and synchronizing data between devices], such as Gigabit Ethernet (GbE), Fibre Channel (FC) or InfiniBand (IB). (Loffink, 2008, p. 9)

Fabric A is dedicated to the Gigabit Ethernet. Although initial server module releases are designed as dual GbE LAN on Motherboard (LOM) controllers on the server module planar, the midplane is enabled to support up to four GbE links per server module on Fabric A. Potential data bandwidth for Fabric A is 4 Gb/s per server module.

Fabrics B and C are identical, fully customizable fabrics that are routed as two sets of four lanes from the mezzanine cards on the server modules to the I/O modules in the rear of the chassis. Supported bandwidth ranges from 1 to 10 Gb/s per lane depending on the fabric type used. Figure 9 illustrates the individual paths of the high-speed I/O architecture.

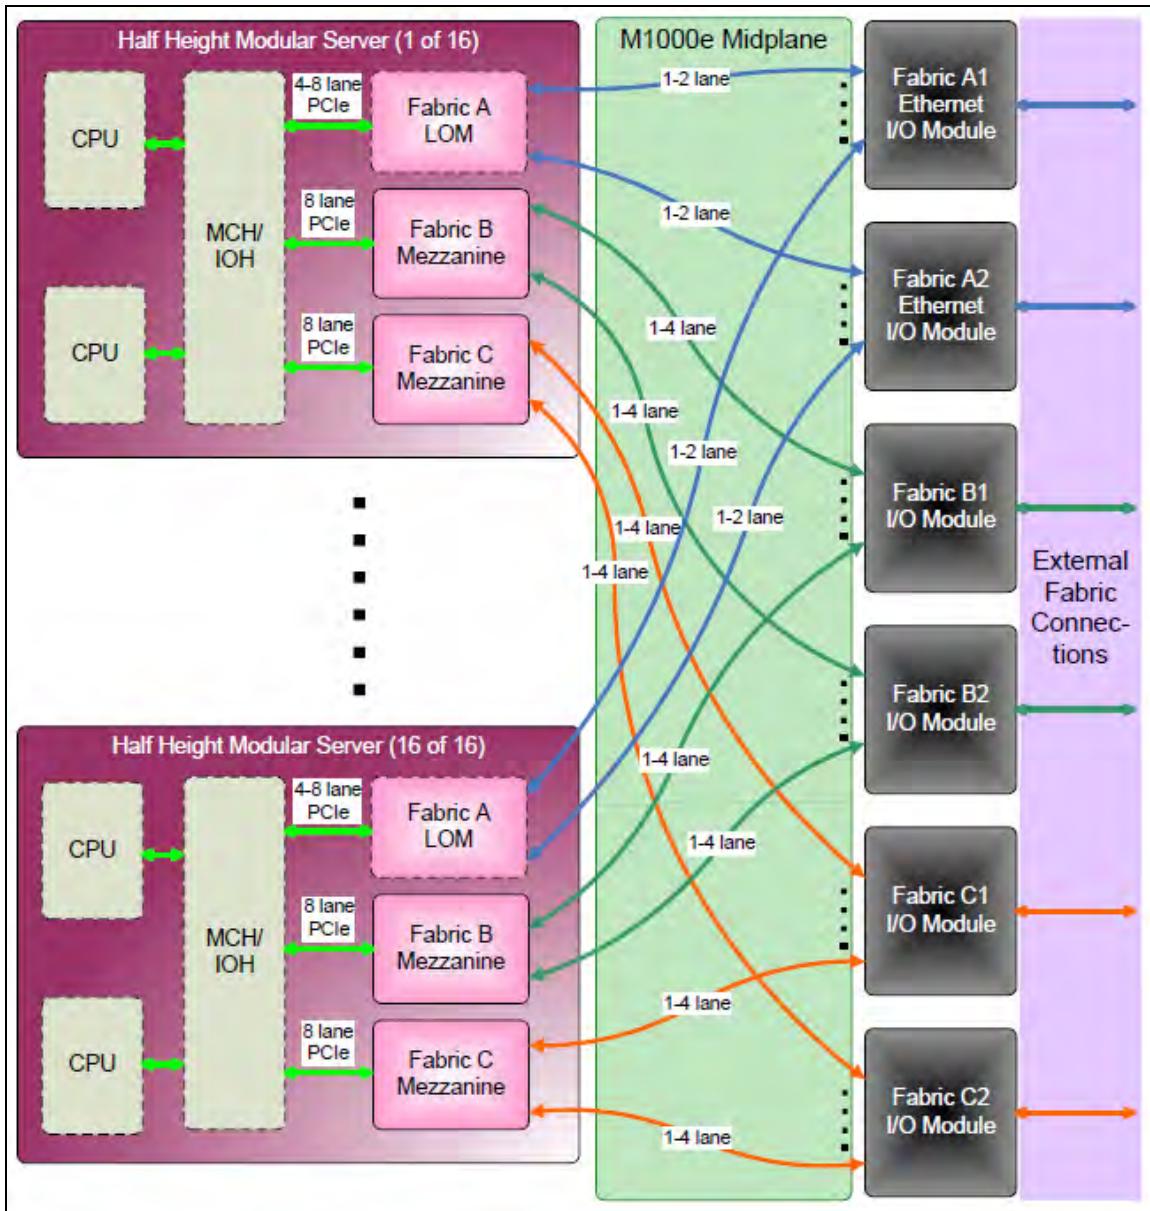


Figure 9. High Speed I/O Architecture (From Loffink, 2008)

In addition to providing maximum scalability and redundancy within its server architecture, Dell also offers options among its line of blade server modules. Selection is dependent upon the unique processing requirements of specific virtual applications. Table 3 details a side-by-side comparison between M605 and M610 blade server modules.

Server Model	M605	M610
<b>Processor</b>	AMD Opteron 2000 series Dual and Quad Core 68W and 95W options	Intel Xeon series Dual and Quad Core 40W, 65W, 80W, and 120W options
<b>Chipset</b>	NVIDIA MCP55	Intel 5000P (Blackford)
<b>Memory Slots</b>	8 DDR2 (667/800 MHz)	8 Fully Buffered DIMMs (667 MHz)
<b>Memory Capacity</b>	32GB (4GB x 8) at launch 64 GB (8GB x 8) planned Q108	32GB (4GB x 8) at launch 64 GB (8GB x 8) planned Q108
<b>LAN on Motherboard (LOM)</b>	2 x GE with hardware TCP/IP Offload Engine and iSCSI Firmware Boot Upgradable to full iSCSI offload via license key	2 x GE with hardware TCP/IP Offload Engine and iSCSI Firmware Boot Upgradable to full iSCSI offload via license key
<b>Fabric Expansion Options</b>	(2) 8-lane PCIe mezzanine cards 1. Dual port GE w/ TOE 2. Dual Port FC4 (Emulex & Qlogic) 3. Dual Port 4x DDR InfiniBand	(2) 8-lane PCIe mezzanine cards 1. Dual port GE w/ TOE 2. Dual Port FC4 (Emulex & Qlogic) 3. Dual Port 4x DDR InfiniBand
<b>Baseboard Management</b>	iDRAC w/ IPMI 2.0 + vMedia + vKVM	iDRAC w/ IPMI 2.0 + vMedia + vKVM
<b>Local Storage Controller Options</b>	SATA (chipset-based: no RAID or hotplug) SAS6/IR (R0/1) CERC6/i (R0/1 w/ Cache)	SATA (chipset-based: no RAID or hotplug) SAS6/IR (R0/1) CERC6/i (R0/1 w/ Cache)
<b>Local Storage Hard Disk Drive (HDD)</b>	(2) 2.5" hot pluggable SAS or SATA	(2) 2.5" hot pluggable SAS or SATA
<b>Video</b>	ATI RN50	ATI RN50
<b>USB</b>	(2) USB 2.0 bootable ports on front panel for floppy/CD/DVD/Memory	(2) USB 2.0 bootable ports on front panel for floppy/CD/DVD/Memory
<b>High-availability (HA) Clustering</b>	Fibre Channel and iSCSI-based clustering options	Fibre Channel and iSCSI-based clustering options

Table 3. Server Module Options (After Loffink, 2008)

#### 4. Power Distribution

Power distribution inside the PowerEdge™ M1000e blade enclosure consists of a 3+3 redundant power supply system. Typical power supply configurations include N+N, N+1, and N+0 redundancy models.

The N+N configuration provides maximum system protection against alternating current (AC) grid loss or power supply failure. If one power grid fails, three power supplies lose their AC source while three power supplies on the other grid remain online, providing sufficient power for the system to continue operating uninterrupted.

The N+1 configuration provides protection only against power supply failures, not grid failures. However, the likelihood of multiple power supplies failing at the same time is remote.

The N+0 configuration provides no intrinsic power protection, therefore any desired system protection must be provided independently at the node or chassis level. Typically this model is used in combination with a High Performance Computing Cluster (HPCC) or other clustered environment where the parallel architecture of the processing nodes across a multiple-system chassis provides adequate redundant functionality. Figure 10 illustrates the interconnectivity of the three distribution models.

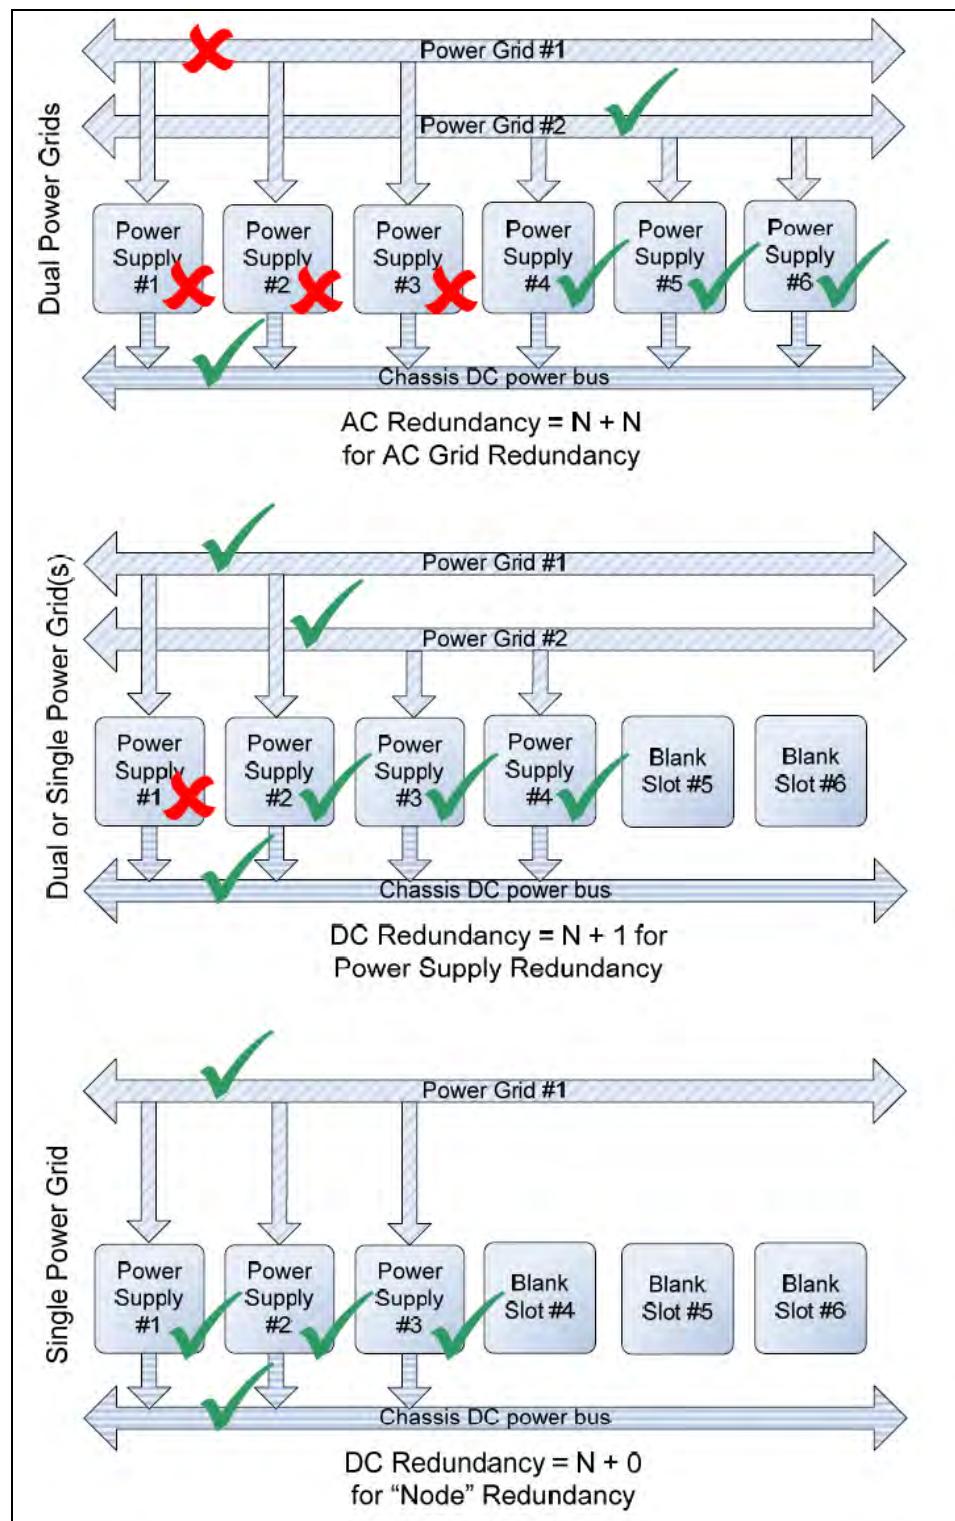


Figure 10. Power Redundancy Modes (From Loffink, 2008)

An improved design to Dell's latest power supply housing includes a fresh air plenum that helps to reduce the air temperature surrounding power supply components. Lower operating temperatures in turn equate to higher power density and higher power efficiency (better than 86% at 20% load) and even higher at heavier loads.

The Power Distribution Unit (PDU) of the PowerEdge™ M1000e blade enclosure is capable of providing stable single-phase output power regardless of whether input power is derived from a single-phase or 3-phase configuration. Table 4 details a side-by-side comparison of the two configurations.

Feature	Single-phase	3-phase
<b>Mounting</b>	1U (horizontal) / zero U (vertical)	
<b>Outlets</b>	3 x C19	
<b>Input Line Voltage</b>	200 to 240 VAC nominal	
<b>Input Line Frequency</b>	47 to 63 Hertz	
<b>Input Line Current</b>	41.4 Amps	24 Amps
<b>Recommended AC service</b>	60 Amps	30 Amps (NA/Japan), 32 Amps (International)
<b>Fixed Input Plug/Cord Rating</b>	IEC-309 60 A Pin & Sleeve Plug	NEMA L15-30P (NA/Japan), IEC 309 4 pole, 4 wire, 380-415 VAC, 32A (International)
<b>Output Rating Voltage</b>	200-240 VAC 60/50 Hz, 1-phase	
<b>Output Rating Current</b>	13.8 Amps	
<b>Output Rating Current (IEC320 C19)</b>	16 Amps	
<b>Power Density</b>	21 Watts per cubic inch	
<b>Power Efficiency</b>	91%, under normal operating conditions	
<b>Circuit Breaker, Over Current Protection</b>	20 Amps, per outlet receptacle	

Table 4. PDU Options (After Loffink, 2008)

## 5. Cooling

The cooling strategy for the PowerEdge™ M1000e blade enclosure supports a low-impedance, high-efficiency design philosophy. Driving lower airflow impedance allows the fan modules to draw air through the system at relatively lower operating

pressures. This produces up to 40% less backpressure than similarly configured designs. The lower backpressure in turn reduces the power consumed by the fan modules while still meeting system airflow requirements.

The low impedance design is coupled with a high-efficiency air moving device. “Efficiency” in this context refers to the work output of the fan as compared to the electrical power required to operate it. The system fan modules operate at efficiencies up to 40% greater than typical axial fan designs. This directly correlates to savings in energy consumption. The enclosure profiles illustrated in Figure 11 and Figure 12 represent the path of airflow with regard to cooling server and I/O modules respectively.

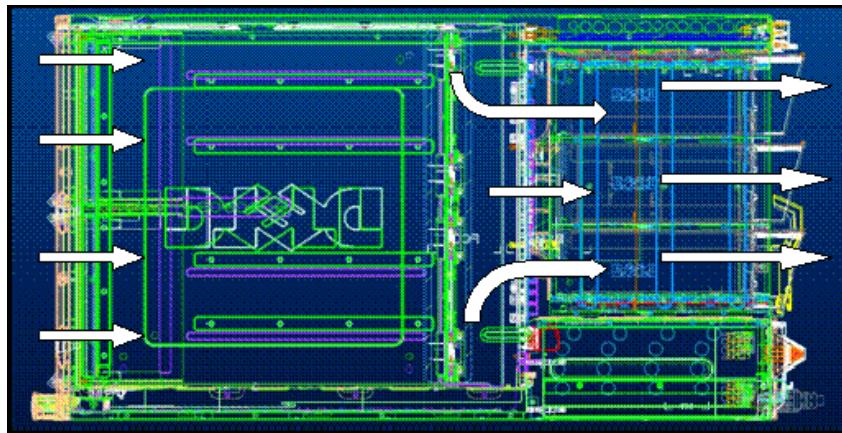


Figure 11. Server Module Cooling Air Profile (From Loffink, 2008)

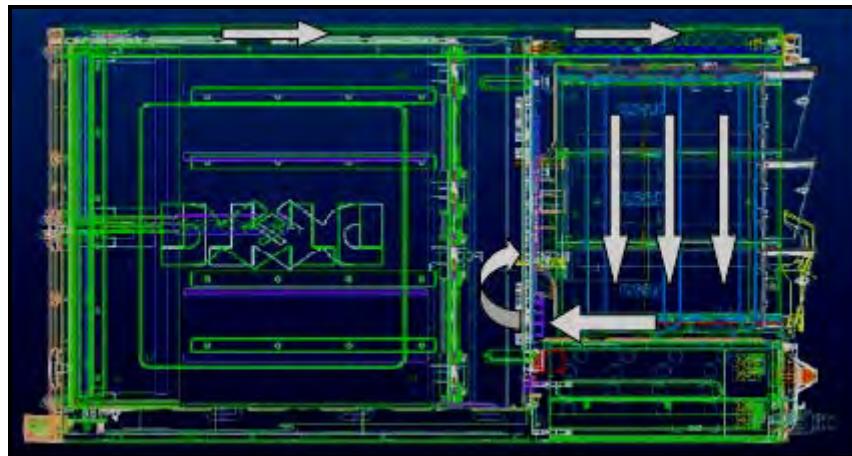


Figure 12. I/O Module Cooling Air Profile (From Loffink, 2008)

## **6. Thermal Control**

The innovation of the physical enclosure design is coupled with a thermal cooling algorithm that allows for dynamic monitoring and adjustment of thermal levels to protect the server and I/O modules.

The iDRAC on each server module calculates the amount of airflow required to reduce the temperature within an individual server module and forwards a request to the CMC. Concurrently, each I/O module (IOM) can send a request to the CMC to increase or decrease cooling to the I/O subsystem. The CMC interprets these requests, and can control the fans as required to maintain module temperatures at optimal levels.

## **7. Remote System Management**

The CMC also provides secure remote management access to the chassis and installed modules. It provides a means for centralized module configuration and direct management of the firmware, firewall traffic, Secure Sockets Layer (SSL)/Secure Shell (SSH), power budget, and remote user access. Remote management also includes monitoring of chassis and module environmental thresholds, such as real-time power consumption, temperature, redundancy, and data consistency.

Each server module's iDRAC comprises the root circuit that integrates the Baseboard Management Controller (BMC) function with hardware support for Virtual Keyboard/Video/Mouse (vKVM) and Virtual Media (vMedia) interfaces for the IP network. Connection to vKVM and vMedia functions is accomplished through the CMC, with encryption available on a per stream basis. The vKVM routes an operator's keyboard, video, and mouse outputs between a physical server and a virtual network console over an IP interface. The vMedia provides for emulation of USB DVD R/W, USB CD R/W, USB Flash Drive, USB ISO image, and USB Floppy over an IP interface. (Loffink, 2008, p. 39)

## **8. Storage Array Network (SAN)**

The Dell iSCSI SAN consists of multiple PS Series storage arrays (arranged in group) that are connected to the IP network and managed as a single system. Each storage array can have up to six network interfaces, arranged in a redundant configuration (3+3) where only three arrays are active at any given time. The other three arrays are in standby mode should one or more network connections fail. Multiple network connections ensure optimal system performance and resource availability.

## **D. SOFTWARE PERFORMANCE**

### **1. Data Traffic Management**

On many networks, it is possible to have an imbalance within the IP network between the devices that send traffic and the devices that receive the traffic. This is often the case in SAN configurations in which many servers (initiators) are communicating with storage devices (targets). If senders transmit data simultaneously, they may exceed the throughput capacity of the receiver. When this occurs, the receiver may drop packets, forcing senders to retransmit the data after a delay. Although this will not result in any loss of data, latency will increase because of the retransmissions, and I/O performance will degrade (Dell, 2008, p. 6).

*Flow Control* can help eliminate this problem. Flow Control is designed to reduce network overhead caused by TCP/IP packet retransmission during periods of heavy data traffic. By allowing the receiver to instruct the sender to briefly slow packet transmission when the receiver is overwhelmed by data packets, the receiver can quickly process its backlog and then resume accepting input. While this functionality might not appear to directly apply to non-TCP/IP internal networks such as the AEGIS computing infrastructure, further studies should be undertaken to determine if this feature could provide additional value to the processing conventions of uplink/downlink traffic between AEGIS ships and missiles in flight.

## **2. Unicast Switching**

A traffic “storm” occurs when a large outpouring of packets creates excessive network traffic that degrades network performance. Many switches have traffic storm control features that prevent ports from being disrupted by broadcast, multicast, or unicast traffic storms on physical interfaces. These features typically work by discarding network packets when the traffic on an interface reaches a percentage of the overall load (usually 80 percent, by default) (Dell, 2008, p. 7).

Since iSCSI unicast traffic typically uses its entire link, the broadcast and multicast storm control feature is probably unnecessary for switches that process only iSCSI traffic. However, it is yet unclear whether (or how) this functionality might apply the AEGIS computing infrastructure. Regardless, this feature is a software configuration item that can be enabled/disabled as required.

## **3. Jumbo Frames**

Ethernet traffic is transported in *frames*. A standard Ethernet frame allows up to 1500 bytes of data to be transferred within a single Ethernet transaction. Jumbo Frames extend Ethernet frames to 9000 bytes to allow more data to be transferred with each Ethernet transaction. Jumbo Frames help reduce the interrupt overhead on the server. Dell® EqualLogic™ PS Series storage arrays support standard Ethernet frames and Jumbo Frames up to 9000 bytes (9018 including the TCP header itself). This is sometimes referred to as the “Alteon standard.” (Dell, 2008, p. 7)

By allowing more data to be transferred with each operation, Jumbo Frames provide an enhanced capacity through faster processing of iSCSI SAN traffic. This would likely prove critical to AEGIS threat engagement, tracking, and prosecution.

## **4. Virtual Local Area Networks (VLANs)**

VLANs use logical connections rather than physical connections, reducing management overhead through greater flexibility. VLANs can be used to accomplish a variety of network settings. Most relevant to AEGIS, and directly applicable to the

previous two sections, is the isolation of iSCSI traffic and selection of specific switches to enable Jumbo Frames. Separation of SAN traffic from other network traffic provides switch-based security enhancements such as port blocking and address filtering.

## 5. Storage Management

All virtualization software enables multiple hosts to share the same physical data repository. VMware® ESX software, chosen for our application, uses a highly-optimized storage stack and Virtual Machine File System (VMFS). VMware maintains that centralized storage of VMs through a VMFS provides more control and flexibility. The VMFS also enables diverse virtualization capabilities such as live migration (VMware calls this capability VMotion), high availability through scheduling of distributed resources, and dynamic clustering.

## 6. Scalability Measures

Evaluating processing scalability is typically defined by *throughput* and *latency*. Throughput is a measure of the amount of data transferred within a given unit of time, expressed in kilobytes per second (KBps) or megabytes per second (MBps). The effectiveness of throughput is dependent on a number of factors such as channel link speed, the volume of unresolved I/O requests, and the system's fundamental caching algorithms.

Latency, on the other hand, is a measure of the time taken to complete a single I/O request, expressed in milliseconds (msec). Latency is dependent on a number of factors such as I/O request size, queue depth (each request must pass through multiple layers of a SAN), physical disk properties, and the system's fundamental caching algorithms.

Figure 13 illustrates that, with the exception of sequential reads, caused by streams arriving from different hosts which must be intermixed at the SAN, there is no drop in aggregate throughput as the number of hosts is scaled up.

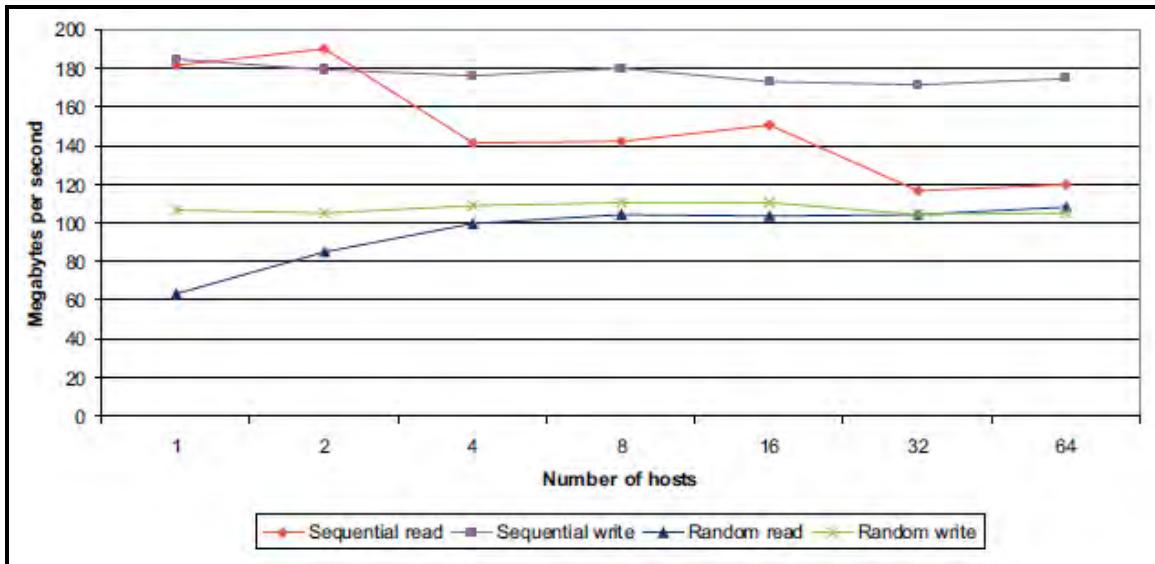


Figure 13. Aggregate I/O Throughput (From VMware, 2008)

Figure 14 illustrates overall latency remains manageable as the number of hosts (32 commands per host) is scaled up. With more than eight hosts, the number of unresolved I/O requests at the SAN's shared storage logical unit (LUN) grows beyond 256 and latencies begin to increase exponentially even though aggregate throughput is steady.

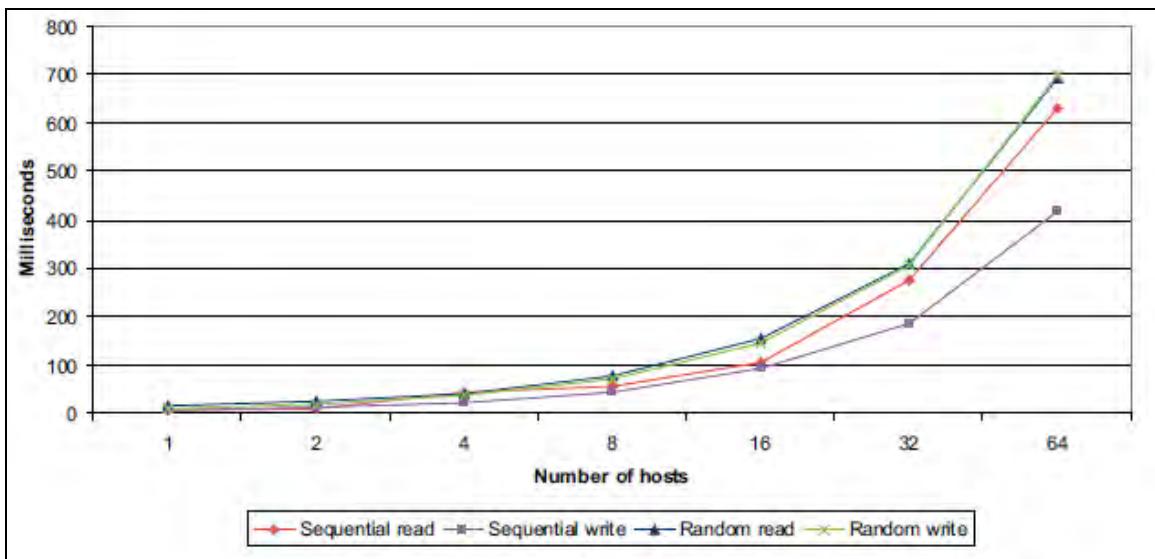


Figure 14. Average I/O Latency (From VMware, 2008)

Figure 15 illustrates the performance of an increasing number of I/O-intensive (45MBps of data) VMs on a single ESX host. When all active paths are routed via the same link (eight VMFS volumes per path), aggregate throughput flattens out after the 2 Gb/s link becomes saturated. When a second link is added between the host and the SAN (four VMFS volumes per link) additional bandwidth is available and the aggregate throughput continues to scale.

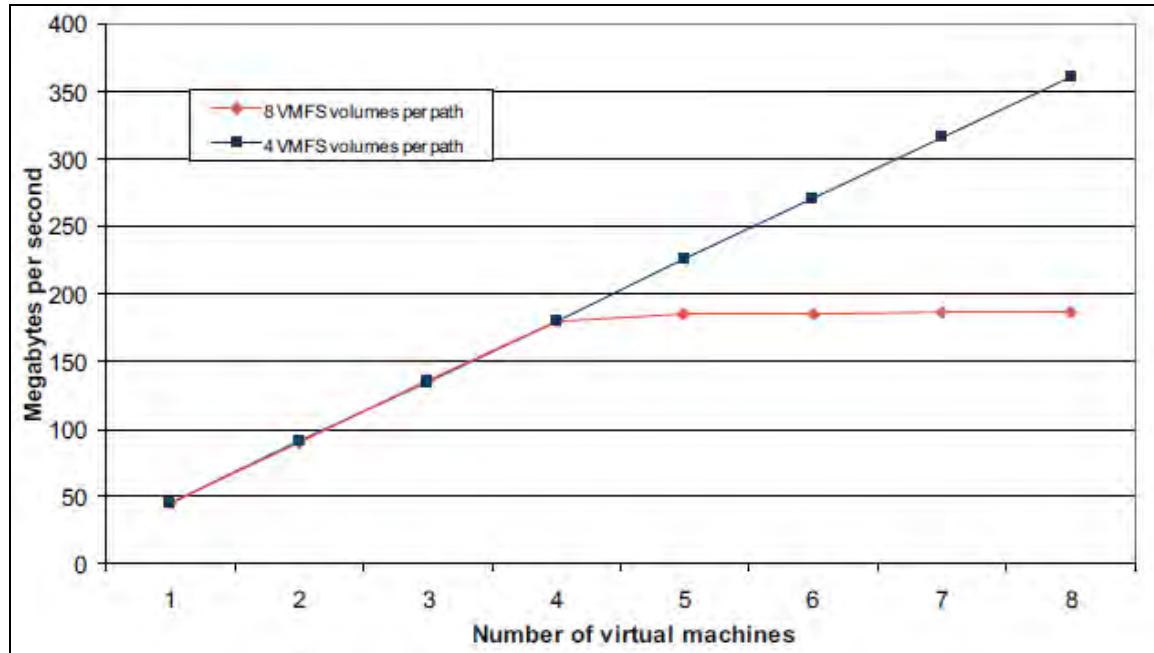


Figure 15. Aggregate Throughput of Multiple I/O-Intensive Virtual Machines (From VMware, 2008)

Figure 16 illustrates the effect on latency during the same test represented in Figure 15. At eight VMFS volumes per link (and beyond four virtual machines) the I/O commands remain in the host bus adapter waiting for the link to become available. Thus, average latency increases as additional VMs are added (a latency of 50 milliseconds is usually a reliable indication that the SAN either does not have enough resources or is not optimally configured to handle its current workload). With two links, however, the aggregate link bandwidth is adequate for the commands to be processed at wire speed.

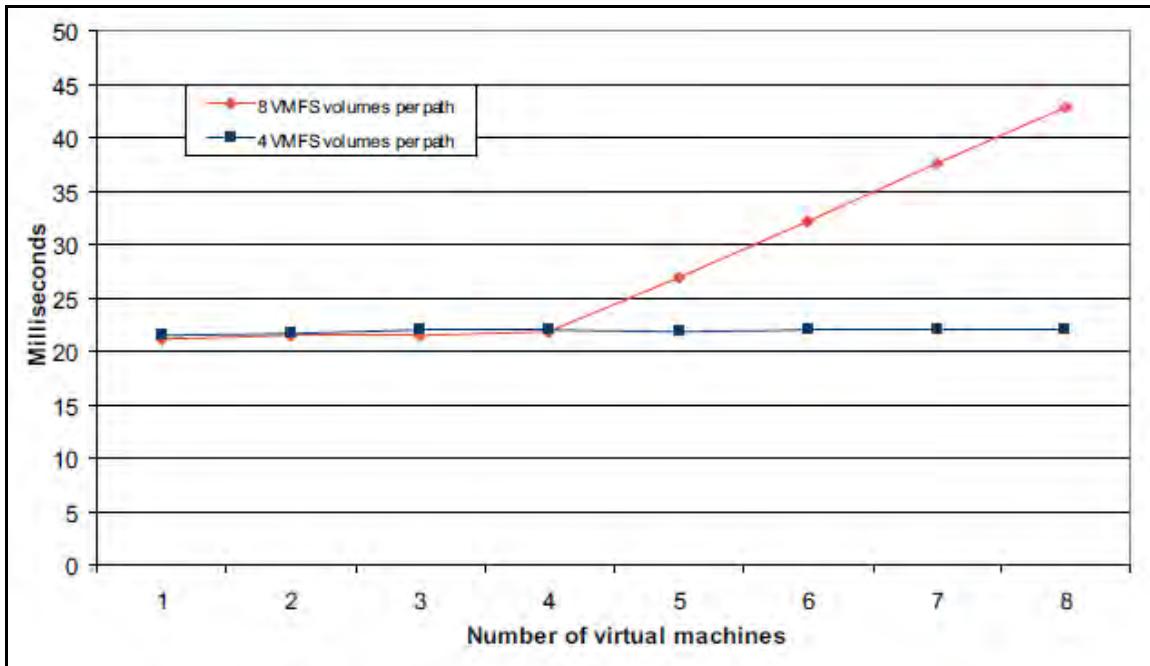


Figure 16. Average Latency of Multiple I/O-Intensive Virtual Machines (From VMware, 2008)

## 7. Application Validity

The modular structure of the Dell PowerEdge™ M1000e server enclosure, combined with M600/605/610 blade servers, provides maximum flexibility through a dynamic I/O, power, cooling, and management architecture. Additionally, optimized power and cooling distribution subsystems maximize component service life and support current and future generations of server and I/O modules.

By offering an I/O bandwidth of up to 40 Gb/s, the Dell server can easily support architectures with multiple 10 Gb Ethernet requirements. All of these features, in addition to demanding a less expensive initial procurement outlay (than comparable monolithic servers), lowers the overall total cost of ownership (TCO).

THIS PAGE INTENTIONALLY LEFT BLANK

## V. CONCLUSION

### A. SUMMARY OF WORK

The AEGIS Weapon System (AWS) and AEGIS Combat System (ACS) modernization efforts will continue to improve the capabilities of cruisers and destroyers against current and future threats. These efforts will also extend the service life and interoperability of components and subsystems.

However, the continued practice of using proprietary equipment and exclusive contracting will only serve to increase costs and reliance on outside assistance to maintain our systems. While Navy technicians are dedicated, intelligent, and highly trained, the vast majority of them do not hold masters degrees in computer science or network operations. Therefore, without access to exclusive tools and knowledge, commanders have no choice but to rely on costly technical assistance from vendors.

Faced with this reality, the problem becomes an economic issue of how to simplify and make less expensive the maintainability of complex systems to the point where qualified operators and technicians can effectively interact with the hardware and software components of systems for which they are responsible.

A related issue that closed source architectures incur is that of diminished manufacturing sources. The average service life of most AEGIS hardware components exceeds eight years. However, the average span of hardware production lines is less than two years. So, the question becomes how to resolve maintainability, usability, and compatibility issues with such short product development timelines?

Typical “old school” solutions include: solicitation/contract with a new vendor well into the equipment’s lifecycle; costly up-front procurement of bulk inventory; desperate mass procurement just prior to end of production; and contractual obligation of mandatory active production lines. While these heavy-handed tactics may have been sufficient in the past, modern budgets are much more constrained and do not allow for such overly expensive reactionary solutions.

Leveraging existing open architecture (OA) technology to improve current computing infrastructures establishes a more solid foundation for reducing costs and providing rapid implementation of future capabilities. Employment of OA systems brings to bear open competition to innovate and achieve improved performance and affordability through the use of modular design and common interface standards.

The ultimate solution, therefore, is a better acquisition strategy through open source architectures that maximize the use of *shared* resources. Plus, virtualization provides the ultimate shared asset scheme by delivering resources on-demand.

Virtualization has the potential to provide significant cost savings in terms of procurement, daily operation, and maintenance. Optimization costs are also reduced through dynamic management of resources such as with load-balancing and processor utilization. Other benefits include streamlined scalability and simplified disaster recovery strategies.

Moreover, the greatest advantage of virtualization, particularly as it pertains to shipboard environments, is the reduced physical footprint of the architecture. As virtual machines (VM) eliminate the need for large computer systems that dominate interior spaces, production of heat and noise is greatly reduced—as well as the energy requirements to power and cool them.

## **B. KEY FINDINGS AND CONTRIBUTIONS**

Various components within the government and military have embraced specific applications of virtualization technology. A prime example is the Consolidated Afloat Networks and Enterprise Services (CANES) program, a current Navy initiative to improve shipboard IT environments through virtualization. At the time of this writing, the Navy (through SPAWAR) has issued nine separate Requests For Information (RFI) with regard to CANES. The latest RFI, issued July 2011, by the Naval Enterprise Networks (NEN) Program Management Office (PMW-205) / Next Generation Enterprise (NGEN) was released for the purpose of seeking industry comments on thin, ultra-thin, and zero client end-user devices.

Interest in thin clients demonstrated the Navy's interest in cost savings through virtualized systems. According to the RFI description, the Department of the Navy (DON) has established a goal of reducing its non-tactical IT budget by 25 percent. This begs the question of why virtualization initiatives should be restricted to non-tactical (or business) applications. The advantages of commodity-based hardware, dynamic processing, scalable storage, reduced physical footprint, and efficient power and cooling requirements would be equally beneficial to tactical computing architectures.

Through construction and analysis of the AEGIS-VM test platform, it was determined that virtualized environments make the most effective use of available resources. However, virtualization can impose a greater load on the storage infrastructure due to increased consolidation levels. Specifically, an I/O command generated within a virtualized environment must pass through extra layers of processing in order to exploit sharing properties. This does not negate all of the useful features of virtualization. However, it is important to understand that potential bottlenecks at various layers exist and that configuration changes may be necessary to attain optimal performance.

### **C. EVALUATION AND FEEDBACK**

The premise behind building the AEGIS-VM test platform was to employ cutting edge open source hardware and software components in the development of a virtualized server platform that would be capable of matching or exceeding the computing performance of current technology employed aboard Navy ships.

The Operational Readiness Test System (ORTS) diagnostic module was specifically targeted to serve as a springboard for comparing performance metrics between the legacy architecture of typical AEGIS configurations and the OA architecture of a Dell® blade server. The ORTS subsystem was chosen because it is the least complex AEGIS subsystem and does not demand the same security considerations as other functional units.

From this foundation, a testing platform comprised solely of open source hardware and software components was designed and constructed. The AEGIS-VM virtual server was designed to be a vehicle from which comparison performance data could be collected and analyzed.

The government software package obtained from SPAWAR was comprised of the ORTS application embedded with the Solaris 8 operating system (OS). Unfortunately, the software was designed to only run on UNIX-based systems and we were therefore unable to use it on our Intel-based system.

It is ironic that the project underlying this thesis, developed in part to demonstrate the utility of using open source products to uncouple compulsory hardware and software associations, should be derailed by a computer program that is only compatible with older OS versions designed to run on specific types of physical processors.

Nevertheless, we can conclude from this study that current virtualization technology can meet the computational requirements of AEGIS. However, OA software is needed before the AEGIS-VM can become a fully functional test platform.

In a very real sense, the inability to run the software in order to prove the value of open source technology is in itself proof. Software applications that are developed with embedded operating systems for specific hardware systems is exactly the type of antiquated architecture that OA and virtualization methodology serves to prevent. In a virtualized environment, operating systems and applications are no longer dependent on each other or on specific hardware configurations.

#### **D. RECOMMENDATION FOR FUTURE WORK**

Due to the time constraints of this thesis research, more detailed analyses and experimentation are required to draw an authoritative conclusion as to the applicability of virtualization technology to the various baselines and subsystems within the AEGIS architecture. Yet there are several directions from which future NPS students can use this thesis as a starting point to advance the knowledge.

A working relationship with the Program Executive Office for Integrated Warfare Systems (PEO IWS 7) at Space and Naval Warfare Systems Command (SPAWAR), Dahlgren, will be a key element. The ORTS program (or other AEGIS application) must be uncoupled from proprietary operating systems and processors.

However, a fully operational test platform is not absolutely necessary to further this research. Practical considerations, such as network security and information assurance, can be evaluated to determine what benefits or adverse effects might be realized from the use of virtualized systems within the AEGIS environment where catastrophic consequences might occur from a security breach.

Other vulnerabilities may also exist. As newer technologies are developed and implemented into legacy systems, greater potential for intrusion exists. Any initiative to build, troubleshoot, or restore a computer program, or to download procedures from an external source needs to account for, and guard against, potential vulnerabilities.

The *flow control* feature of the PowerEdge™ servers is designed to reduce network overhead caused by TCP/IP packet retransmission during periods of heavy data traffic. While this functionality might not appear to directly apply to non-TCP/IP internal networks such as the AEGIS computing infrastructure, further studies should be undertaken to determine if this feature could provide additional value to the processing conventions of uplink/downlink traffic between AEGIS ships and missiles in flight.

The Consolidated Afloat Networks and Enterprise Services (CANES) program, a current Navy IT initiative to improve shipboard IT environments through virtualization, may be another inroad to expand the Navy's interest in virtualized systems beyond non-tactic applications. The goal of CANES is to improve shipboard networks by consolidating five legacy systems currently in use aboard Navy ships. This seems to be a good fit with the goal to move AEGIS away from its open/closed hybrid architecture.

At the time of this writing, the Navy (through SPAWAR) had just issued the ninth Request For Information (RFI) with regard to CANES. The latest RFI, issued July 2011, by the Naval Enterprise Networks (NEN) Program Management Office (PMW-205) / Next Generation Enterprise (NGEN) was focused on thin, ultra-thin, and zero client end-

user devices. An interesting (and positive) footnote to the RFI is that it specifically stated that proprietary information would *not* be accepted.

The value of this line of research cannot be overstated. As benchmarks are established, and lessons learned gained, cost savings—through a steady transition of computing systems from sole source development to genuine open architectures—will be particularly critical as budgetary constraints continue to tighten. AEGIS, in particular, will be at the forefront as limited Navy resources are stretched thin in the face of evolving national security threats.

## LIST OF REFERENCES

Adler, J. R., & Ahart, J. L. (2007). *AEGIS Platforms: Using KVA Analysis to Assess Open Architecture in Sustaining Engineering* (Master's Thesis). Naval Postgraduate School, Monterey, CA.

*AEGIS Combat System (ACS)*. Retrieved April 6, 2010 from Global Security Web site: <http://www.globalsecurity.org/military/systems/ship/systems/aegis.htm>

*AEGIS Open Architecture (OA)*. Retrieved April 6, 2010 from Global Security Web site: <http://www.globalsecurity.org/military/systems/ship/systems/aegis-oa.htm>

Brazet, M. (1994). AEGIS ORTS—the first and future ultimate integrated diagnostics system. *Aerospace and Electronic Systems Magazine, IEEE*, 9(2), 40–45. Current version released August 2002 in *IEEE Xplore*. doi: 10.1109/62.260044.

Corrin, Amber (2011, January 19). *Navy's next-gen shipboard IT passes critical milestone*. Retrieved from Defense Systems Web site: <http://defensesystems.com/articles/2011/01/19/navy-canis-milestone-b-approval.aspx>

Dell, Inc. (June 2008). *PS Series Array Network Performance Guidelines*. Technical Report. Retrieved August 25, 2010 from Dell EqualLogic Web site: <http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229>

Ewing, Philip (2010, June 28). *U.S. AEGIS Radars' Readiness Plunges*. Retrieved from Defense News Web site: <http://www.defensenews.com/story.php?i=4688283>

Filz, Warren (2009, October 14). *AEGIS Combat System / AEGIS Weapon System overview brief for Naval Postgraduate School*. PowerPoint presentation.

Housel, T. & Mun, J. (2007). *AEGIS Platforms: The Potential Impact of Open Architecture in Sustaining Engineering* (NPS-AM-07-053). Retrieved from Naval Postgraduate School Web site: [http://www.acquisitionresearch.net/\\_files/FY2007/NPS-AM-07-053.pdf](http://www.acquisitionresearch.net/_files/FY2007/NPS-AM-07-053.pdf)

Loffink, J. (2008). *Dell PowerEdge M1000e Modular Enclosure Architecture*. Dell Enterprise White Paper. Retrieved September 21, 2010 from Dell Web site: [http://www.dell.com/downloads/global/products/pedge/en/pedge\\_m1000e\\_white\\_paper.pdf](http://www.dell.com/downloads/global/products/pedge/en/pedge_m1000e_white_paper.pdf)

Mell, P. & Grance, T. (2009). *NIST working definition of cloud computing* (draft). Retrieved June 3, 2010 from National Institute of Standard and Technology Web site: [http://info.apps.gov/sites/default/files/NIST\\_Cloud\\_Definition.doc](http://info.apps.gov/sites/default/files/NIST_Cloud_Definition.doc)

Moore, G. (1965). Cramming more components onto integrated circuits. *Electronics*, 114–117. Reprinted January 1998 in *Proceedings of the IEEE*, 86(1), 82–85. doi: S 0018-9219(98)00753-1.

Richfield, Paul (2010, November 17). *Navy sets course to better link shipboard networks*. Retrieved from Defense Systems Web site:  
<http://defensesystems.com/articles/2010/11/17/c4isr-1-navy-shipboard-network-upgrades.aspx>

Sun Microsystems (2009, June 16). *End of Service Life Status for Solaris Operating System*. Retrieved from Sun Web site:  
[http://www.sun.com/service/eosl/eosl\\_Solaris.html](http://www.sun.com/service/eosl/eosl_Solaris.html)

Thompson, S. (2008). *Extending Open Architecture to the Physical Layer* (Master's Thesis). Naval Postgraduate School, Monterey, CA.

Tiglao, L. (2010). *Virtual Machine Modules for use by DoD C4I Support Centers* (Master's Thesis). Naval Postgraduate School, Monterey, CA.

Uchytil, J. (2006). *Assessing the Operational Value of Situational Awareness for AEGIS and Ship Self Defense System (SSDS) Platforms through the Application of the Knowledge Value Added (KVA) Methodology* (Master's Thesis). Naval Postgraduate School, Monterey, CA.

VMware, Inc. (2008). *Scalable Storage Performance: VMware® ESX 3.5. Performance Study*. Retrieved August 25, 2010 from VMware Web site:  
[http://www.vmware.com/files/pdf/scalable\\_storage\\_performance.pdf](http://www.vmware.com/files/pdf/scalable_storage_performance.pdf)

Williams, J. (2003). *AEGIS Combat Systems Where Are We Now?*. Paper presented at ASNE Day 2003 meeting of the American Society of Naval Engineers, Arlington, VA.

## **INITIAL DISTRIBUTION LIST**

1. Defense Technical Information Center  
Ft. Belvoir, Virginia
2. Dudley Knox Library  
Naval Postgraduate School  
Monterey, California
3. Dr. Dan Boger  
Naval Postgraduate School  
Monterey, California
4. Dr. Man-Tak Shing  
Naval Postgraduate School  
Monterey, California
5. Professor Albert Barreto III  
Naval Postgraduate School  
Monterey, California
6. Commander Kurt Rothenhaus  
C4I Program Executive Office  
PMW 160 Tactical Networks  
San Diego, California